

ATTORNEY'S DOCKET NO. RIN-001
PATENTS

UNITED STATES PATENT APPLICATION
OF

NICHOLAS A. J. MILLINGTON

FOR

SYSTEM AND METHOD FOR SYNCHRONIZING OPERATIONS AMONG A PLURALITY OF
INDEPENDENTLY CLOCKED DIGITAL DATA PROCESSING DEVICES

Certificate of Express Mailing

Express Mail Mailing Label No. ET 492 359 690 US

Date of Deposit April 1, 2004

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office To Addressee" Service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D. C. 20231.

By Richard A. Jordan

Richard A. Jordan

INCORPORATION BY REFERENCE

This application claims the benefit of Provisional Patent Application Serial No. 60/490,768, filed on July 28, 2003, entitled "Method for Synchronizing Audio Playback Between Multiple Networked Devices," assigned to the assignee of the present application, incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates generally to the field of digital data processing devices, and more particularly to systems and methods for synchronizing operations among a plurality of independently-clocked digital data processing devices. The invention is embodied in a system for synchronizing operations among a plurality of devices, in relation to information that is provided by a common source. One embodiment of the invention enables synchronizing of audio playback as among two or more audio playback devices that receive audio information from a common information source, or channel.

More generally, the invention relates to the field of arrangements that synchronize output generated by a number of output generators, including audio output, video output, combinations of audio and video, as well as other types of output as will be appreciated by those skilled in the art, provided by a common channel. Generally, the invention will find utility in connection with any type of information for which synchrony among independently-clocked devices is desired.

BACKGROUND OF THE INVENTION

There are a number of circumstances under which it is desirable to maintain synchrony of operations among a plurality of independently-clocked digital data processing devices in relation to, for example, information that is provided thereto by a common source. For example, systems are

1 being developed in which one audio information source can distribute audio information in digital
2 form to a number of audio playback devices for playback. The audio playback devices receive the
3 digital information and convert it to analog form for playback. The audio playback devices may be
4 located in the same room or they may be distributed in different rooms in a residence such as a house
5 or an apartment, in different offices in an office building, or the like. For example, in a system
6 installed in a residence, one audio playback device may be located in a living room, while another
7 audio playback device is be located in a kitchen, and yet other audio playback devices may be
8 located in various bedrooms of a house. In such an arrangement, the audio information that is
9 distributed to various audio playback devices may relate to the same audio program, or the
10 information may relate to different audio programs. If the audio information source provides audio
11 information relating to the same audio program to two or more audio playback devices at the same
12 time, the audio playback devices will generally contemporaneously play the same program. For
13 example, if the audio information source provides audio information to audio playback devices
14 located in the living room and kitchen in a house at the same time, they will generally
15 contemporaneously play the same program.

16 One problem that can arise is to ensure that, if two or more audio playback devices are
17 contemporaneously attempting to play back the same audio program, they do so simultaneously.
18 Small differences in the audio playback devices' start times and/or playback speeds can be perceived
19 by a listener as an echo effect, and larger differences can be very annoying. Differences can arise
20 because for a number of reasons, including delays in the transfer of audio information over the
21 network. Such delays can differ as among the various audio playback devices for a variety of
22 reasons, including where they are connected into the network, message traffic and other reasons as
23 will be apparent to those skilled in the art.

24 Another problem arises from the following. When an audio playback device converts the
25 digital audio information from digital to analog form, it does so using a clock that provides timing
26 information. Generally, the audio playback devices that are being developed have independent
27 clocks, and, if they are not clocking at precisely the same rate, the audio playback provided by the
28 various devices can get out of synchronization.

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for synchronizing operations among a number of digital data processing devices that are regulated by independent clocking devices. Generally, the invention will find utility in connection with any type of information for which synchrony among devices connected to a network is desired. The invention is described in connection with a plurality of audio playback devices that receive digital audio information that is to be played back in synchrony, but it will be appreciated that the invention can find usefulness in connection with any kind of information for which coordination among devices that have independent clocking devices would find utility.

In brief summary, the invention provides, in one aspect, a system for maintaining synchrony of operations among a plurality of devices that have independent clocking arrangements. The system includes a task distribution device that distributes tasks to a synchrony group comprising a plurality of devices that are to perform the tasks distributed by the task distribution device in synchrony. The task distribution device distributes each task to the members of the synchrony group over a network. Each task is associated with a time stamp that indicates a time, relative to a clock maintained by the task distribution device, at which the members of the synchrony group are to execute the task. Each member of the synchrony group periodically obtains from the task distribution device an indication of the current time indicated by its clock, determines a time differential between the task distribution device's clock and its respective clock and determines therefrom a time at which, according to its respective clock, the time stamp indicates that it is to execute the task.

In one embodiment, the tasks that are distributed include audio information for an audio track that is to be played by all of the devices comprising the synchrony group synchronously. The audio track is divided into a series of frames, each of which is associated with a time stamp indicating the time, relative to the clock maintained by an audio information channel device, which, in that embodiment, serves as the task distribution device, at which the members of the synchrony group are to play the respective frame. Each member of the synchrony group, using a very accurate

1 protocol, periodically obtains the time indicated by the audio information channel device, and
2 determines a differential between the time as indicated by its local clock and the audio information
3 channel device's clock. The member uses the differential and the time as indicated by the time stamp
4 to determine the time, relative to its local clock, at which it is to play the respective frame. The
5 members of the synchrony group do this for all of the frames, and accordingly are able to play the
6 frames in synchrony.

7 **BRIEF DESCRIPTION OF THE DRAWINGS**

8 This invention is pointed out with particularity in the appended claims. The above and
9 further advantages of this invention may be better understood by referring to the following
10 description taken in conjunction with the accompanying drawings, in which:

11 FIG. 1 schematically depicts an illustrative networked audio system, constructed in
12 accordance with the invention;

13 FIG. 2 schematically depicts a functional block diagram of a synchrony group utilizing a
14 plurality of zone players formed within the networked audio system depicted in FIG. 1;

15 FIG. 2A schematically depicts two synchrony groups, illustrating how a member of one
16 synchrony group can provide audio information to the members of another synchrony group;

17 FIG. 3 depicts an functional block diagram of a zone player for use in the networked audio
18 system depicted in FIG. 1; and

19 FIG. 4 is useful in understanding an digital audio information framing methodology useful
20 in the network audio system depicted in FIG. 1.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 depicts an illustrative network audio system 10 constructed in accordance with the invention. With reference to FIG. 1, the network audio system 10 includes a plurality of zone players 11(1) through 11(N) (generally identified by reference numeral 11(n)) interconnected by a local network 12, all of which operate under control of one or more user interface modules generally identified by reference numeral 13. One or more of the zone players 11(n) may also be connected to one or more audio information sources, which will generally be identified herein by reference numeral 14(n)(s), and/or one or more audio reproduction devices, which will generally be identified by reference numeral 15(n)(r). In the reference numeral 14(n)(s), index "n" refers to the index "n" of the zone player 11(n) to which the audio information source is connected, and the index "s" ($s=1, \dots, S_n$) refers to the "s-th" audio information source connected to that "n-th" zone player 11(n). Thus, if, for example, a zone player 11(n) is connected to four audio information sources 14(n)(1) through 14(n)(4), the audio information sources may be generally identified by reference numeral 14(n)(s), with $S_n=4$. It will be appreciated that the number of audio information sources S_n may vary as among the various zone players 11(n), and some zone players may not have any audio information sources connected thereto. Similarly, in the reference numeral 15(n)(r), index "n" refers to the index "n" of the zone player 11(n) to which the audio reproduction device is connected, and the index "r" ($r=1, \dots, R_n$) refers to the "r-th" audio information source connected to that "n-th" zone player 11(n). In addition to the audio information sources 14(n)(s), the network audio system 10 may include one or more audio information sources 16(1) through 16(M) connected through appropriate network interface devices (not separately shown) to the local network 12. Furthermore, the local network may include one or more network interface devices (also not separately shown) that are configured to connect the local network 12 to other networks, including a wide area network such as the Internet, the public switched telephony network (PSTN) or other networks as will be apparent to those skilled in the art, over which connections to audio information sources may be established.

1 The zone players 11(n) associated with system 10 may be distributed throughout an
2 establishment such as residence, an office complex, a hotel, a conference hall, an amphitheater or
3 auditorium, or other types of establishments as will be apparent to those skilled in the art or the like.
4 For example, if the zone players 11(n) and their associated audio information source(s) and/or audio
5 reproduction device(s) are distributed throughout a residence, one, such as zone player 11(1) and its
6 associated audio information source(s) and audio reproduction device(s) may be located in a living
7 room, another may be located in a kitchen, another may be located in a dining room, and yet others
8 may be located in respective bedrooms, to selectively provide entertainment in those rooms. On the
9 other hand, if the zone players 11(n) and their associated audio information source(s) and/or audio
10 reproduction device(s) are distributed throughout an office complex, one may, for example, be
11 provided in each office to selectively provide entertainment to the employees in the respective
12 offices. Similarly, if the zone players 11(n) and associated audio information source(s) and/or audio
13 reproduction device(s) are used in a hotel, they may be distributed throughout the rooms to provide
14 entertainment to the guests. Similar arrangements may be used with zone players 11(n) and
15 associated audio information source(s) and/or audio reproduction device(s) used in an amphitheater
16 or auditorium. Other arrangements in other types of environments will be apparent to those skilled
17 in the art. In each case, the zone players 11(n) can be used to selectively provide entertainment in
18 the respective locations, as will be described below.

19 The audio information sources 14(n)(s) and 16(m) may be any of a number of types of
20 conventional sources of audio information, including, for example, compact disc ("CD") players,
21 AM and/or FM radio receivers, analog or digital tape cassette players, analog record turntables and
22 the like. In addition, the audio information sources 14(n)(s) and 16(m) may comprise digital audio
23 files stored locally on, for example, personal computers (PCs), personal digital assistants (PDAs),
24 or similar devices capable of storing digital information in volatile or non-volatile form. As noted
25 above, the local network 12 may also have an interface (not shown) to a wide area network, over
26 which the network audio system 10 can obtain audio information. Moreover, one or more of the
27 audio information sources 14(n)(s) may also comprise an interface to a wide area network such as
28 the Internet, the public switched telephony network (PSTN) or any other source of audio

1 information. In addition, one or more of the audio information sources 14(n)(s) and 16(m) may
2 comprise interfaces to radio services delivered over, for example, satellite. Audio information
3 obtained over the wide area network may comprise, for example, streaming digital audio information
4 such as Internet radio, digital audio files stored on servers, and other types of audio information and
5 sources as will be appreciated by those skilled in the art. Other arrangements and other types of
6 audio information sources will be apparent to those skilled in the art.

7 Generally, the audio information sources 14(n)(s) and 16(m) provide audio information
8 associated with audio programs to the zone players for playback. A zone player that receives audio
9 information from an audio information source 14(n)(s) that is connected thereto can provide
10 playback and/or forward the audio information, along with playback timing information, over the
11 local network 12 to other zone players for playback. Similarly, each audio information source
12 16(m) that is not directly connected to a zone player can transmit audio information over the network
13 12 to any zone player 11(n) for playback. In addition, as will be explained in detail below, the
14 respective zone player 11(n) can transmit the audio information that it receives either from an audio
15 information source 14(n)(s) connected thereto, or from an audio information source 16(m), to
16 selected ones of the other zone players 11(n'), 11(n''),... (n not equal to n', n'',...) for playback by
17 those other zone players. The other zone players 11(n'), 11(n''),... to which the zone player 11(n)
18 transmits the audio information for playback may be selected by a user using the user interface
19 module 13. In that operation, the zone player 11(n) will transmit the audio information to the
20 selected zone players 11(n'), 11(n''),... over the network 12. As will be described below in greater
21 detail, the zone players 11(n), 11(n'), 11(n''),... operate such that the zone players 11(n'), 11(n''),...
22 synchronize their playback of the audio program with the playback by the zone player 11(n), so that
23 the zone players 11(n), 11(n'), 11(n'') provide the same audio program at the same time.

24 Users, using user interface module 13, may also enable different groupings or sets of zone
25 players to provide audio playback of different audio programs synchronously. For example, a user,
26 using a user interface module 13, may enable zone players 11(1) and 11(2) to play one audio
27 program, audio information for which may be provided by, for example, one audio information
28 source 14(1)(1). The same or a different user may, using the same or a different user interface

1 module 13, enable zone players 11(4) and 11(5) to contemporaneously play another audio program,
2 audio information for which may be provided by a second audio information source, such as audio
3 information source 14(5)(2). Further, a user may enable zone player 11(3) to contemporaneously
4 play yet another audio program, audio information for which may be provided by yet another audio
5 information source, such as audio information source 16(1). As yet another possibility, a user may
6 contemporaneously enable zone player 11(1) to provide audio information from an audio
7 information source connected thereto, such as audio information source 14(1)(2), to another zone
8 player, such as zone player 11(6) for playback.

9 In the following, the term "synchrony group" will be used to refer to a set of one or more
10 zone players that are to play the same audio program synchronously. Thus, in the above example,
11 zone players 11(1) and 11(2) comprise one synchrony group, zone player 11(3) comprises a second
12 synchrony group, zone players 11(4) and 11(5) comprise a third synchrony group, and zone player
13 11(6) comprises yet a fourth synchrony group. Thus, while zone players 11(1) and 11(2) are playing
14 the same audio program, they will play the audio program synchronously. Similarly, while zone
15 players 11(4) and 11(5) are playing the same audio program, they will play the audio program
16 synchronously. On the other hand, zone players that are playing different audio programs may do
17 so with unrelated timings. That is, for example, the timing with which zone players 11(1) and 11(2)
18 play their audio program may have no relationship to the timing with which zone player 11(3), zone
19 players 11(4) and 11(5), and zone player 11(6) play their audio programs. It will be appreciated that,
20 since "synchrony group" is used to refer to sets of zone players that are playing the same audio
21 program synchronously, zone player 11(1) will not be part of zone player 11(6)'s synchrony group,
22 even though zone player 11(1) is providing the audio information for the audio program to zone
23 player 11(6).

24 In the network audio system 10, the synchrony groups are not fixed. Users can enable them
25 to be established and modified dynamically. Continuing with the above example, a user may enable
26 the zone player 11(1) to begin providing playback of the audio program provided thereto by audio
27 information source 14(1)(1), and subsequently enable zone player 11(2) to join the synchrony group.
28 Similarly, a user may enable the zone player 11(5) to begin providing playback of the audio program

1 provided thereto by audio information source 14(5)(2), and subsequently enable zone player 11(4)
2 to join that synchrony group. In addition, a user may enable a zone player to leave a synchrony
3 group and possibly join another synchrony group. For example, a user may enable the zone player
4 11(2) to leave the synchrony group with zone player 11(1), and join the synchrony group with zone
5 player 11(6). As another possibility, the user may enable the zone player 11(1) to leave the
6 synchrony group with zone player 11(2) and join the synchrony group with zone player 11(6). In
7 connection with the last possibility, the zone player 11(1) can continue providing audio information
8 from the audio information source 14(1)(1) to the zone player 11(2) for playback thereby.

9 A user, using the user interface module 13, can enable a zone player 11(n) that is currently
10 not a member of a synchrony group to join a synchrony group, after which it will be enabled to play
11 the audio program that is currently being played by that synchrony group. Similarly, a user, also
12 using the user interface module 13, can enable a zone player 11(n) that is currently a member of one
13 synchrony group, to disengage from that synchrony group and join another synchrony group, after
14 which that zone player will be playing the audio program associated with the other synchrony group.
15 For example, if a zone player 11(6) is currently not a member of any synchrony group, it, under
16 control of the user interface module 13, can become a member of a synchrony group, after which it
17 will play the audio program being played by the other members of the synchrony group, in
18 synchrony with the other members of the synchrony group. In becoming a member of the synchrony
19 group, zone player 11(6) can notify the zone player that is the master device for the synchrony group
20 that it wishes to become a member of its synchrony group, after which that zone player will also
21 transmit audio information associated with the audio program, as well as timing information, to the
22 zone player 11(6). As the zone player 11(6) receives the audio information and the timing
23 information from the master device, it will play the audio information with the timing indicated by
24 the timing information, which will enable the zone player 11(6) to play the audio program in
25 synchrony with the other zone player(s) in the synchrony group.

26 Similarly, if a user, using the user interface module 13, enables a zone player 11(n)
27 associated with a synchrony group to disengage from that synchrony group, and if the zone player
28 11(n) is not the master device of the synchrony group, the zone player 11(n) can notify the master

1 device, after which the master device can terminate transmission of the audio information and timing
2 information to the zone player 11(n). If the user also enables the zone player 11(n) to begin playing
3 another audio program using audio information from an audio information source 14(n)(s) connected
4 thereto, it will acquire the audio information from the audio information source 14(n)(s) and initiate
5 playback thereof. If the user enables another zone player 11(n') to join the synchrony group
6 associated with zone player 11(n), operations in connection therewith can proceed as described
7 immediately above.

8 As yet another possibility, if a user, using the user interface module 13, enables a zone player
9 11(n) associated with a synchrony group to disengage from that synchrony group and join another
10 synchrony group, and if the zone player is not the master device of the synchrony group from which
11 it is disengaging, the zone player 11(n) can notify the master device of the synchrony group from
12 which it is disengaging, after which that zone player will terminate transmission of audio
13 information and timing information to the zone player 11(n) that is disengaging.
14 Contemporaneously, the zone player 11(n) can notify the master device of the synchrony group that
15 it (that is, zone player 11(n)) is joining, after which the master device can begin transmission of
16 audio information and timing information to that zone player 11(n). The zone player 11(n) can
17 thereafter begin playback of the audio program defined by the audio information, in accordance with
18 the timing information so that the zone player 11(n) will play the audio program in synchrony with
19 the master device.

20 As yet another possibility, a user, using the user interface module 13, may enable a zone
21 player 11(n) that is not associated with a synchrony group, to begin playing an audio program using
22 audio information provided to it by an audio information source 14(n)(s) connected thereto. In that
23 case, the user, also using the user interface module 13 or a user interface device that is specific to
24 the audio information source 14(n)(s), can enable the audio information source 14(n)(s) to provide
25 audio information to the zone player 11(n). After the zone player 11(n) has begun playback, or
26 contemporaneously therewith, the user, using the user interface module 13, can enable other zone
27 players 11(n'), 11(n''),... to join zone player 11(n)'s synchrony group and enable that zone player

11(n) to transmit audio information and timing information thereto as described above, to facilitate synchronous playback of the audio program by the other zone players 11(n'), 11(n'')....

A user can use the user interface module 13 to control other aspects of the network audio system 10, including but not limited to the selection of the audio information source 14(n)(s) that a particular zone player 11(n) is to utilize, the volume of the audio playback, and so forth. In addition, a user may use the user interface module 13 to turn audio information source(s) 14(n)(s) on and off and to enable them to provide audio information to the respective zone players 11(n).

Operations performed by the various devices associated with a synchrony group will be described in connection with FIG. 2, which schematically depicts a functional block diagram of a synchrony group in the network audio system 10 described above in connection with FIG. 1. With reference to FIG. 2, a synchrony group 20 includes a master device 21 and zero or more slave devices 22(1) through 22 (G) (generally identified by reference numeral 22(g)), all of which synchronously play an audio program provided by an audio information channel device 23. Each of the master device 21, slave devices 22(g) and audio information channel device 23 utilizes a zone player 11(n) depicted in FIG. 1, although it will be clear from the description below that a zone player may be utilized both for the audio information channel device for the synchrony group 20, and the master device 21 or a slave device 22(g) of the synchrony group 20. As will be described below in more detail, the audio information channel device 23 obtains the audio information for the audio program from an audio information source, adds playback timing information, and transmits the combined audio and playback timing information to the master device 21 and slave devices 22(g) over the network 12 for playback. The playback timing information that is provided with the audio information, together with clock timing information provided by the audio information channel device 23 to the various devices 21 and 22(g) as will be described below, enables the master device 21 and slave devices 22(g) of the synchrony group 20 to play the audio information simultaneously.

The master device 21 and the slave devices 22(g) receive the audio and playback timing information, as well as the clock timing information, that are provided by the audio information channel device 23, and play back the audio program defined by the audio information. The master

1 device 21 is also the member of the synchrony group 20 that communicates with the user interface
2 module 13 and that controls the operations of the slave devices 22(g) in the synchrony group 20.
3 In addition, the master device 21 controls the operations of the audio information channel device 23
4 that provides the audio and playback timing information for the synchrony group 20. Generally, the
5 initial master device 21 for the synchrony group will be the first zone player 11(n) that a user wishes
6 to play an audio program. However, as will be described below, the zone player 11(n) that operates
7 as the master device 21 can be migrated from one zone player 11(n) to another zone player 11(n'),
8 which preferably will be a zone player that is currently operating as a slave device 22(g) in the
9 synchrony group.

10 In addition, under certain circumstances, as will be described below, the zone player 11(n)
11 that operates as the audio information channel device 23 can be migrated from one zone player to
12 another zone player, which also will preferably will be a zone player that is currently operating as
13 a member of the synchrony group 20. It will be appreciated that the zone player that operates as the
14 master device 21 can be migrated to another zone player independently of the migration of the audio
15 information channel device 23. For example, if one zone player 11(n) is operating as both the master
16 device 21 and the audio information channel device 23 for a synchrony group 20, the master device
17 21 can be migrated to another zone player 11(n') while the zone player 11(n) is still operating as the
18 audio information channel device 23. Similarly, if one zone player 11(n) is operating as both the
19 master device 21 and the audio information channel device 23 for a synchrony group 20, the audio
20 information channel device 23 can be migrated to another zone player 11(n') while the zone player
21 11(n) is still operating as the master device 21. In addition, if one zone player 11(n) is operating as
22 both the master device 21 and the audio information channel device 23 for a synchrony group 20,
23 the master device 21 can be migrated to another zone player 11(n') and the audio information
24 channel device can be migrated to a third zone player 11(n").

25 The master device 21 receives control information from the user interface module 13 for
26 controlling the synchrony group 20 and provides status information indicating the operational status
27 of the synchrony group to the user interface module 13. Generally, the control information from the
28 user interface module 13 enables the master device 21 to, in turn, enable the audio information

1 channel device 23 to provide audio and playback timing information to the synchrony group to
2 enable the devices 21 and 22(g) that are members of the synchrony group 20 to play the audio
3 program synchronously. In addition, the control information from the user interface module 13
4 enables the master device 21 to, in turn, enable other zone players to join the synchrony group as
5 slave devices 22(g) and to enable slave devices 22(g) to disengage from the synchrony group.
6 Control information from the user interface module 13 can also enable the zone player 11(n) that is
7 currently operating as the master device 21 to disengage from the synchrony group, but prior to
8 doing so that zone player will enable the master device 21 to transfer from that zone player 11(n) to
9 another zone player 11(n'), preferably to a zone player 11(n') that is currently a slave device 22(g) in
10 the synchrony group 20. The control information from the user interface module 13 can also enable
11 the master device 21 to adjust its playback volume and to enable individual ones of the various slave
12 devices 22(g) to adjust their playback volumes. In addition, the control information from the user
13 interface module 13 can enable the synchrony group 20 to terminate playing of a current track of the
14 audio program and skip to the next track, and to re-order tracks in a play list of tracks defining the
15 audio program that is to be played by the synchrony group 20.

16 The status information that the master device 21 may provide to the user interface module
17 13 can include such information as a name or other identifier for the track of the audio work that is
18 currently being played, the names or other identifiers for upcoming tracks, the identifier of the zone
19 player 11(n) that is currently operating as the master device 21, and identifiers of the zone players
20 that are currently operating as slave devices 22(g). In one embodiment, the user interface module
21 13 includes a display (not separately shown) that can display the status information to the user.

22 It will be appreciated that the zone player 11(n) that is operating as the audio information
23 channel device 23 for one synchrony group may also comprise the master device 21 or any of the
24 slave devices 22(g) in another synchrony group. This may occur if, for example, the audio
25 information source that is to provide the audio information that is to be played by the one synchrony
26 group is connected to a zone player also being utilized as the master device or a slave device for the
27 other synchrony group. This will be schematically illustrated below in connection with FIG. 2A.
28 Since, as noted above, the zone player 11(n) that is operating as the audio information channel

1 device 23 for the synchrony group 20 may also be operating as a master device 21 or slave device
2 22(g) for another synchrony group, it can also be connected to one or more audio reproduction
3 devices 15(n)(r), although that is not depicted in FIG. 2. Since the master device 21 and slave
4 devices 22(g) are all to provide playback of the audio program, they will be connected to respective
5 audio reproduction devices 15(n)(r). Furthermore, it will be appreciated that one or more of the
6 zone players 11(n) that operate as the master device 21 and slave devices 22(g) in synchrony group
7 20 may also operate as an audio information channel device for that synchrony group or for another
8 synchrony group and so they may be connected to one or more audio information sources 14(n)(s),
9 although that is also not depicted in FIG. 2. In addition, it will be appreciated that a zone player
10 11(n) can also operate as a audio information channel device 23 for multiple synchrony groups.

11 If the audio information channel device 23 does not utilize the same zone player as the master
12 device 21, the master device 21 controls the audio information channel device by exchanging control
13 information over the network 12 with the audio information channel device 23. The control
14 information is represented in FIG. 2 by the arrow labeled CHAN_DEV_CTRL_INFO. The control
15 information that the master device 21 provides to the audio information channel device 23 will
16 generally depend on the nature of the audio information source that is to provide the audio
17 information for the audio program that is to be played and the operation to be enabled by the control
18 information. If, for example, the audio information source is a conventional compact disc, tape, or
19 record player, broadcast radio receiver, or the like, which is connected to a zone player 11(n), the
20 master device 21 may merely enable the zone player serving as the audio information channel device
21 23 to receive the audio information for the program from the audio information source. It will be
22 appreciated that, if the audio information is not in digital form, the audio information channel device
23 23 will convert it to digital form and provide the digitized audio information, along with the
24 playback timing information, to the master device 21 and slave devices 22(g).

25 On the other hand, if the audio information source is, for example, a digital data storage
26 device, such as may be on a personal computer or similar device, the master device 21 can provide
27 a play list to the audio information channel device 23 that identifies one or more files containing the
28 audio information for the audio program. In that case, the audio information channel device 23 can

1 retrieve the files from the digital data storage device and provide them, along with the playback
2 timing information, to the master device 21 and the slave devices 22(g). It will be appreciated that,
3 in this case, the audio information source may be directly connected to the audio information channel
4 device 23, as, for example, an audio information source 14(n)(s), or it may comprise an audio
5 information source 16(m) connected to the network 12. As a further alternative, if the audio
6 information source is a source available over the wide area network, the master device 21 can
7 provide a play list comprising a list of web addresses identifying the files containing the audio
8 information for the audio program that is to be played, and in that connection the audio information
9 channel device 23 can initiate a retrieval of the files over the wide area network. As yet another
10 alternative, if the audio information source is a source of streaming audio received over the wide area
11 network, the master device 21 can provide a network address from which the streaming audio can
12 be received. Other arrangements by which the master device 21 can control the audio information
13 channel device 23 will be apparent to those skilled in the art.

14 The master device 21 can also provide control information to the synchrony group's audio
15 information channel device 23 to enable a migration from one zone player 11(n) to another zone
16 player 11(n'). This may occur if, for example, the audio information source is one of audio
17 information sources 16 or a source accessible over the wide area network via the network 12. The
18 master device 21 can enable migration of the audio information channel device 23 for several
19 reasons, including, for example, to reduce the loading of the zone player 11(n), to improve latency
20 of message transmission in the network 12, and other reasons as will be appreciated by those skilled
21 in the art.

22 As noted above, the audio information channel device 23 provides audio and playback timing
23 information for the synchrony group to enable the master device 21 and slave devices 22(g) to play
24 the audio program synchronously. Details of the audio and playback timing information will be
25 described in detail below in connection with FIGS. 3 and 4, but, in brief, the audio information
26 channel device 23 transmits the audio and playback timing information in messages over the network
27 12 using a multi-cast message transmission methodology. In that methodology, the audio
28 information channel device 23 will transmit the audio and playback timing information in a series

1 of messages, with each message being received by all of the zone players 11(n) comprising the
2 synchrony group 20, that is, by the master device 21 and the slave devices 22(g). Each of the
3 messages includes a multi-cast address, which the master device 21 and slave devices 22(g) will
4 monitor and, when they detect a message with that address, they will receive and use the contents
5 of the message. The audio and playback timing information is represented in FIG. 2 by the arrow
6 labeled "AUD+PBTIME_INF0," which has a single tail, representing a source for the information
7 at the audio information channel device 23, and multiple arrowheads representing the destinations
8 of the information, with one arrowhead extending to the master device 21 and other arrowheads
9 extending to each of the slave devices 22(g) in the synchrony group 20. The audio information
10 channel device 23 may make use of any convenient multi-cast message transmission methodology
11 in transmitting the audio and playback timing information to the synchrony group 20. As will be
12 described in detail in connection with FIG. 4, the audio and playback timing information is in the
13 form of a series of frames, with each frame having a time stamp. The time stamp indicates a time,
14 relative to the time indicated by a clock maintained by the audio information channel device 23, at
15 which the frame is to be played. Depending on the size or sizes of the messages used in the selected
16 multi-cast message transmission methodology and the size or sizes of the frames, a message may
17 contain one frame, or multiple frames, or, alternatively, a frame may extend across several messages.

18 The audio information channel device 23 also provides clock time information to the master
19 device 21 and each of the slave devices 22(g) individually over network 12 using a highly accurate
20 clock time information transmission methodology. The distribution of the clock time information
21 is represented in FIG. 2 by the arrows labeled "AICD_CLK_INF (M)" (in the case of the clock time
22 information provided to the master device 21) and "AICD_CLK_INF (S₁)" through
23 "AICD_CLK_INF (S_G)" (in the case of audio information channel device clock information provided
24 to the slave devices 22(g)). In one embodiment, the master device 21 and slave devices 22(g) make
25 use of the well-known SNTP (Simple Network Time Protocol) to obtain current clock time
26 information from the audio information channel device 23. The SNTP makes use of a unicast
27 message transfer methodology, in which one device, such as the audio information channel device
28 23, provides clock time information to a specific other device, such as the master device 21 or a

1 slave device 22(g), using the other device's network, or unicast, address. Each of the master device
2 21 and slave devices 22(g) will periodically initiate SNTP transactions with the audio information
3 channel device 23 to obtain the clock time information from the audio information channel device
4 23. As will be described below in more detail, the master device 21 and each slave device 22(g)
5 make use of the clock time information to determine the time differential between the time indicated
6 by the audio information channel device's clock and the time indicated by its respective clock, and
7 use that time differential value, along with the playback time information associated with the audio
8 information and the respective device's local time as indicated by its clock to determine when the
9 various frames are to be played. This enables the master device 21 and the slave devices 22(g) in
10 the synchrony group 20 to play the respective frames simultaneously.

11 As noted above, the control information provided by the user to the master device 21 through
12 the user interface module 13 can also enable the master device 21 to, in turn, enable another zone
13 player 11(n') to join the synchrony group as a new slave device 22(g). In that operation, the user
14 interface module 13 will provide control information, including the identification of the zone player
15 11(n') that is to join the synchrony group to the master device 21. After it receives the identification
16 of the zone player 11(n') that is to join the synchrony group, the master device 21 will exchange
17 control information, which is represented in FIG. 2 by the arrows labeled SLV_DEV_CTRL_INF
18 (S_1) through SLV_DEV_CTRL_INF (S_G) group slave control information, over the network 12 with
19 the zone player 11(n') that is identified in the control information from the user interface module 13.
20 The control information that the master device 21 provides to the new zone player 11(n') includes
21 the network address of the zone player 11(n) that is operating as the audio information channel
22 device 23 for the synchrony group, as well as the multi-cast address that the audio information
23 channel device 23 is using to broadcast the audio and playback timing information over the network.
24 The zone player that is to operate as the new slave device 22(g') uses the multi-cast address to begin
25 receiving the multi-cast messages that contain the audio information for the audio program being
26 played by the synchrony group.

27 It will be appreciated that, if the zone player 11(n) that is operating as the master device 21
28 for the synchrony group 20 is also operating the audio information channel device 23, and if there

1 are no slave devices 22(g) in the synchrony group 20, the audio information channel device 23 may
2 not be transmitting audio and playback timing information over the network. In that case, if the new
3 slave device 22(g') is the first slave device in the synchrony group, the zone player 11(n) that is
4 operating as both the master device 21 and audio information channel device 23, can begin
5 transmitting the audio and playback timing information over the network 12 when the slave device
6 22(g') is added to the synchrony group 20. The zone player 11(n) can maintain a count of the
7 number of slave devices 22(g) in the synchrony group 20 as they join and disengage, and, if the
8 number drops to zero, it can stop transmitting the audio and playback timing information over the
9 network 12 to reduce the message traffic over the network 12.

10 The new slave device 22(g') added to the synchrony group 20 uses the network address of
11 the audio information channel device 23 for several purposes. In particular, the new slave device
12 22(g') will, like the master device 21 (assuming the zone player 11(n) operating as the master device
13 21 is not also the audio information channel device 23), engage in SNTP transactions with the audio
14 information channel device 23 to obtain the clock timing information from the audio information
15 channel device 23. In addition, the new slave device 22(g') can notify the audio information channel
16 device 23 that it is a new slave device 22(g') for the synchrony group 20 and provide the audio
17 information channel device 23 with its network address. As will be described below, in one
18 embodiment, particularly in connection with audio information obtained from a source, such as a
19 digital data storage device, which can provide audio information at a rate that is faster than the rate
20 at which it will be played, the audio information channel device 23 will buffer audio and timing
21 information and broadcast it over the network 12 to the synchrony group 20 generally at a rate at
22 which it is provided by the source. Accordingly, when a new slave device 22(g') joins the synchrony
23 group 20, the playback timing information may indicate that the audio information that is currently
24 being broadcast by the audio information channel device 23 using the multi-cast methodology is to
25 be played back some time in the future. To reduce the delay with which the new slave device 22(g')
26 will begin playback, the audio information channel device 23 can also retransmit previously
27 transmitted audio and timing information that it had buffered to the new slave device 22(g') using
28 the unicast network address of the slave device 22(g').

1 The master device 21 can also use the slave device control information exchanged with the
2 slave devices 22(g) for other purposes. For example, the master device 21 can use the slave device
3 control information to initiate a migration of the master from its zone player 11(n) to another zone
4 player 11(n'). This may occur for any of a number of reasons, including, for example, that the master
5 device 21 is terminating playback by it of the audio program and is leaving the synchrony group 20,
6 but one or more of the other devices in the synchrony group is to continue playing the audio
7 program. The master device 21 may also want to initiate a migration if it is overloaded, which can
8 occur if, for example, the zone player 11(n) that is the master device 21 for its synchrony group is
9 also operating as an audio information channel device 23 for another synchrony group.

10 The user can also use the user interface module 13 to adjust playback volume by the
11 individual zone players 11(n) comprising the synchrony group. In that operation, the user interface
12 module 13 provides information identifying the particular device whose volume is to be adjusted,
13 and the level at which the volume is to be set to the master device 21. If the device whose volume
14 is to be adjusted is the master device 21, the master device 21 can adjust its volume according to the
15 information that it receives from the user interface module 13. On the other hand, if the device
16 whose volume is to be adjusted is a slave device 22(g), the master device 21 can provide group slave
17 control information to the respective slave device 22(g), to enable it to adjust its volume.

18 The user can also use the user interface module 13 to enable a synchrony group 20 to cancel
19 playing of the track in an audio program that is currently being played, and to proceed immediately
20 to the next track. This may occur, for example, if the tracks for the program is in the form of a series
21 of digital audio information files, and the user wishes to cancel playback of the track that is defined
22 by one of the files. In that case, when the master device 21 receives the command to cancel playback
23 of the current track, it will provide channel device control information to the audio information
24 channel device 23 so indicating. In response, the audio information channel device 23 inserts control
25 information into the audio and playback timing information, which will be referred to as a
26 "resynchronize" command. In addition, the audio information channel device 23 will begin
27 transmitting audio information for the next track, with timing information to enable it to be played

1 immediately. The resynchronize command can also enable playback of a track to be cancelled
2 before it has been played. Details of these operations will be described below.

3 As noted above, there may be multiple synchrony groups in the network audio system 10,
4 and further that, for example, a zone player 11(n) may operate both as a master device 21 or a slave
5 device 22(g) in one synchrony group, and as the audio information channel device 23 providing
6 audio and playback timing information and clock timing information for another synchrony group.
7 An illustrative arrangement of this will be described in connection with FIG. 2A. With reference
8 to FIG. 2A, that FIG. depicts elements of two synchrony groups, identified by reference numerals
9 20(1) and 20(2), respectively. For clarity, FIG. 2A does not show a number of elements, the
10 presence of which would be evident from FIGS. 1 and 2 as described above. For example, FIG. 2A
11 does not depict the audio information sources from which audio information is obtained for the
12 synchrony groups or the audio reproduction devices that are used to produce sound for the master
13 and slave devices, which are depicted in both FIGS. 1 and 2. In addition, FIG. 2A does not depict
14 arrows that represent control information provided by the respective master devices to the slave
15 devices in the respective synchrony groups, or to the audio information channel devices that provide
16 audio and timing information for the respective synchrony groups, which are depicted in FIG. 2. In
17 addition, FIG. 2A does not depict the arrows that represent the clock timing information provided
18 by the audio information channel devices to the respective members of the respective synchrony
19 groups, which are also depicted in FIG. 2. As will be noted below, however, FIG. 2A does depict
20 arrows representing the audio and playback timing information provided by the respective audio
21 information channel devices for the respective synchrony groups 20(1), 20(2), to the master and
22 slave devices comprising the respective synchrony groups 20(1), 20(2).

23 Each synchrony group 20(1), 20(2) comprises elements of a number of zone players. A
24 functional block diagram of a zone player will be described below in connection with FIG. 3.
25 Synchrony group 20(1) includes a master device 21(1) and "K" slave devices 22(1)(1) through
26 22(K)(1) (the index "1" in reference numeral 21(1) and the last index in reference numeral 22(1)(1)
27 through 21(K)(1) corresponds to the index of the synchrony group 20(1) to which they belong)
28 utilize zone players 11(1) and 11(K+1) respectively. Similarly, synchrony group 20(2) includes a

1 master device 21(2) and "L" slave devices 22(1)(2) through 22(L)(2) that utilize zone players
2 11(K+2) through 11(K+L+2). In the illustrative arrangement depicted in FIG. 2A, both synchrony
3 groups 20(1) and 20(2) are controlled by the user interface module 13, which can provide control
4 information to, and receive status information from, the master devices 21(1) and 21(2)
5 independently. It will be appreciated that separate user interface modules may be provided to
6 provide control information to, and receive status information from, the respective master devices
7 21(1), 21(2).

8 As noted above, the slave device 22(1)(2) in synchrony group 20(2) utilizes zone player
9 11(K+3). In the illustrative arrangement depicted in FIG. 2A, the audio information channel device
10 23(1) that provides audio and playback timing information to the master and slave devices 21(1),
11 22(1)(1),..., 22(K)(1) of synchrony group 20(1) also utilizes zone player 11(K+3). As noted above,
12 this may occur if, for example, the audio information source that is to provide audio information to
13 be played by the synchrony group 20(1) is connected to the zone player 11(K+3). Thus, when the
14 master device 21(1) of synchrony group 20(1) exchanges channel device control information with
15 the audio information channel device 23(1), it is effectively exchanging channel device control
16 information with the zone player 11(K+3). Similarly, when the master and slave devices 21(1),
17 22(1)(1),..., 22(K)(1) of synchrony group 20(1) receive audio and playback timing information, as
18 well as clock timing information, from the audio information channel device 23(1), they are
19 effectively receiving the information from the zone player 11(K+3). FIG. 2A depicts a multi-headed
20 arrow representing audio and playback timing information transmitted by the zone player 11(K+3),
21 as audio information channel device 23(1), to the master and slave devices 21(1),
22 22(1)(1),..., 11(K)(1) comprising synchrony group 20(1).

23 On the other hand, in the illustrative arrangement depicted in FIG. 2A, the synchrony group
24 20(2) utilizes a zone player 11(K+L+3) as its audio information channel device 23(2). As with
25 synchrony group 20(1), when the master device 21(2) of synchrony group 20(2) exchanges channel
26 device control information with the audio information channel device 23(2), it is effectively
27 exchanging channel device control information with the zone player 11(K+L+3). Similarly, when
28 the master and slave devices 21(2), 22(1)(2),..., 22(L)(2) of synchrony group 20(2) receive audio and

1 playback timing information, as well as clock timing information, from the audio information
2 channel device 23(2), they are effectively receiving the information from the zone player
3 11(K+L+3). FIG. 2A depicts a multi-headed arrow representing audio and playback timing
4 information transmitted by the zone player 11(K+3) as audio information channel device 23(2) to
5 the master and slave devices 21(2), 22(1)(2),...,22(L)(2) comprising synchrony group 20(2).

6 In the illustrative arrangement depicted in FIG. 2A, zone player 11(K+L+3), which is the
7 audio information channel device 23(2) for synchrony group 20(2), is not shown as being either a
8 master or a slave device in another synchrony group. However, it will be appreciated that zone
9 player 11(K+L+3) could also be utilized as the master device or a slave device for another synchrony
10 group. Indeed, it will be appreciated that the zone player that is utilized as the audio information
11 channel device for synchrony group 20(2) may also be a zone player that is utilized as the master
12 device 21(1) or a slave device 22(1)(1),..., 22(K)(1) in the synchrony group 20(1).

13 A zone player 11(n) that is utilized as a member of one synchrony group may also be utilized
14 as the audio information channel device for another synchrony group if the audio information source
15 that is to supply the audio information that is to be played by the other synchrony group is connected
16 to that zone player 11(n). A zone player 11(n) may also be utilized as the audio information channel
17 device for the other synchrony group if, for example, the audio information source is an audio
18 information source 16(m) (FIG. 1) that is connected to the network 12 or an audio information
19 source that is available over a wide area network such as the Internet. The latter may occur if, for
20 example, the zone player 11(n) has sufficient processing power to operate as the audio information
21 channel device and it is in an optimal location in the network 12, relative to the zone players
22 comprising the other synchrony group (that is the synchrony group for which it is operating as audio
23 information channel device) for providing the audio and playback timing information to the
24 members of the other synchrony group. Other circumstances under which the zone player 11(n) that
25 is utilized as a member of one synchrony group may also be utilized as the audio information
26 channel device for another synchrony group will be apparent to those skilled in the art.

1 As was noted above, the master device 21 for a synchrony group 20 may be migrated from
2 one zone player 11(n) to another zone player 11(n'). As was further noted above, the audio
3 information channel device 23 for a synchrony group 20 may be migrated from one zone player
4 11(n) to another zone player 11(n'). It will be appreciated that the latter may occur if, for example,
5 the audio information source that provides the audio program for the synchrony group is not
6 connected to the zone player 11(n) that is operating as the audio information channel device 23, but
7 instead is one of the audio information sources 16(m) connected to the network 12 or a source
8 available over a wide area network such as the Internet. Operations performed during a migration
9 of an audio information channel device 23 from one zone player 11(n) to another zone player 11(n')
10 will generally depend on the nature of the audio information that is being channeled by the audio
11 information channel device 23. For example, if the audio information source provides streaming
12 audio, the zone player 11(n) that is currently operating as the audio information channel device 23
13 for the synchrony group 20, can provide the following information to the other zone player 11(n')
14 that is to become the audio information channel device 23 for the synchrony group 20:

15 (a) the identification of the source of the streaming audio information,

16 (b) the time stamp associated with the frame that the zone player 11(n) is currently forming,

17 and

18 (c) the identifications of the zone players that are operating as the master device 21 and slave
19 devices 22(g) comprising the synchrony group 20.

20 After the zone player 11(n') receives the information from the zone player 11(n), it will begin
21 receiving the streaming audio from the streaming audio information source identified by the zone
22 player 11(n), assemble the streaming audio information into frames, associate each frame with a time
23 stamp, and transmit the resulting audio and playback timing information over the network 12. The
24 zone player 11(n') will perform these operations in the same manner as described above, except that,
25 instead of using the time indicated by its digital to analog converter clock 34 directly in generating
26 the time stamps for the frames, the initial time stamp will be related to the value of the time stamp
27 that is provided by the zone player 11(n) (reference item (b) above), with the rate at which the time

1 stamps are incremented corresponding to the rate at which its (that is, the zone player 11(n')s) clock
2 increments. In addition, the zone player 11(n') will notify the zone players that are operating as the
3 master device 21 and slave devices 22(g) of the synchrony group 20 that it is the new audio
4 information channel device 23 for the synchrony group 20, and provide the multi-cast address that
5 it will be using to multi-cast the audio and playback timing information, as well as its unicast
6 network address. After the members of the synchrony group 20 receive the notification from the
7 zone player 11(n') indicating that it is the new audio information channel device 23 for the synchrony
8 group 20, they will receive the audio and playback timing information from the zone player 11(n')
9 instead of the zone player 11(n), using the multi-cast address provided by the zone player 11(n').
10 In addition, they can utilize the zone player 11(n')s unicast network address to obtain current time
11 information therefrom. It will be appreciated that the zone player 11(n') will determine its current
12 time in relation to the time stamp that is provided by the zone player 11(n) (reference item (b) above)
13 or the current time information that it received from the zone player 11(n) using the SNTP protocol
14 as described above.

15 Generally similar operations can be performed in connection with migrating the audio
16 information channel device from one zone player 11(n) to another zone player 11(n') if the audio
17 information is from one or more audio information files, such as may be the case if the audio
18 information comprises MP3 or WAV files that are available from sources such as sources 16(m)
19 connected to the network 12 or over from sources available over a wide area network such as the
20 Internet, except for differences to accommodate the fact that the audio information is in files. In that
21 case, the zone player 11(n) that is currently operating as the audio information channel device 23 for
22 the synchrony group 20 can provide the following information to the zone player 11(n') that is to
23 become the audio information channel device 23 for the synchrony group 20:

24 (d) a list of the audio information files containing the audio information that is to be played;

25 (e) the identification of the file for which the zone player 11(n) is currently providing audio
26 and playback timing information, along with the offset into the file for which the current item of

1 audio and playback timing information is being generated and the time stamp that the zone player
2 11(n) is associating with that frame, and

3 (f) the identifications of the zone players that comprise the master device 21 and slave
4 devices 22(g) comprising the synchrony group 20.

5 After the zone player 11(n') receives the information from the zone player 11(n), it will begin
6 retrieving audio information from the file identified in item (e), starting at the identified offset. In
7 addition, the zone player 11(n') can assemble the retrieved audio information into frames, associate
8 each frame with a time stamp and transmit the resulting audio and playback timing information over
9 the network 12. The zone player 11(n') will perform these operations in the same manner as
10 described above, except that, instead of using the time indicated by its digital to analog converter
11 clock 34 directly in generating the time stamps for the frames, the value of the initial time stamp will
12 be related to the time stamp that is provided by the zone player 11(n) (reference item (e) above), with
13 the rate at which the time stamps are incremented corresponding to the rate at which its (that is, the
14 zone player 11(n')'s) clock increments. In addition, the zone player 11(n') will notify the zone
15 players that are operating as the master device 21 and slave devices 22(g) of the synchrony group
16 20 that it is the new audio information channel device 23 for the synchrony group 20, and provide
17 the multi-cast address that it will be using to multi-cast the audio and playback timing information,
18 as well as its unicast network address. After the members of the synchrony group 20 receive the
19 notification from the zone player 11(n') indicating that it is the new audio information channel device
20 23 for the synchrony group 20, they will receive the audio and playback timing information from
21 the zone player 11(n') instead of the zone player 11(n), using the multi-cast address provided by the
22 zone player 11(n'). In addition, they can utilize the zone player 11(n')'s unicast network address to
23 obtain current time information therefrom. It will be appreciated that the zone player 11(n') will
24 determine its current time in relation to the time stamp that is provided by the zone player 11(n)
25 (reference item (b) above) or the current time information that it received from the zone player 11(n)
26 using the SNTP protocol as described above. The zone player 11(n') will process successive audio
27 information files in the list that it receives from the zone player 11(n) (reference item (d)).

1 Operations performed by the zone players 11(n) and 11(n') in connection with migration of
2 the audio information channel device 23 for other types of audio information will be apparent to
3 those skilled in the art. In any case, preferably, the zone player 11(n) will continue operating as an
4 audio information channel device 23 for the synchrony group 20 for at least a brief time after it
5 notifies the zone player 11(n') that it is to become audio information channel device for the
6 synchrony group, so that the zone player 11(n') will have time to notify the zone players in the
7 synchrony group 20 that it is the new audio information channel device 23 for the synchrony group.

8 Before proceeding further in describing operations performed by the network audio system
9 10, it would be helpful to provide a detailed description of a zone player 11(n) constructed in
10 accordance with the invention. FIG. 3 depicts a functional block diagram of a zone player 11(n)
11 constructed in accordance with the invention. All of the zone players in the network audio system
12 10 may have similar construction. With reference to FIG. 3, the zone player 11(n) includes an audio
13 information source interface 30, an audio information buffer 31, a playback scheduler 32, a digital
14 to analog converter 33, an audio amplifier 35, an audio reproduction device interface 36, a network
15 communications manager 40, and a network interface 41, all of which operate under the control of
16 a control module 42. The zone player 11(n) also has a device clock 43 that provides timing signals
17 that control the general operations of the zone player 11(n). In addition, the zone player 11(n)
18 includes a user interface module interface 44 that can receive control signals from the user interface
19 module 13 (FIGS. 1 and 2) for controlling operations of the zone player 11(n), and provide status
20 information to the user interface module 13.

21 Generally, the audio information buffer 31 buffers audio information, in digital form, along
22 with playback timing information. If the zone player 11(n) is operating as the audio information
23 channel device 23 (FIG. 2) for a synchrony group 20, the information that is buffered in the audio
24 information buffer 31 will include the audio and playback timing information that will be provided
25 to the devices 21 and 22(g) in the synchrony group 20. If the zone player 11(n) is operating as the
26 master device 21 or a slave device 22(g) for a synchrony group, the information that is buffered in
27 the audio information buffer 31 will include the audio and playback timing information that the zone
28 player 11(n) is to play.

1 The audio information buffer 31 can receive audio and playback timing information from two
2 sources, namely, the audio information source interface 30 and the network communications
3 manager 40. In particular, if the zone player 11(n) is operating as the audio information channel
4 device 23 for a synchrony group 20, and if the audio information source is a source 14(n)(s)
5 connected to the zone player 11(n), the audio information buffer 31 will receive and buffer audio and
6 playback timing information from the audio information source interface 30. On the other hand, if
7 the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group
8 20, and if the audio information source is a source 16(m) connected to the network 12, or a source
9 available over the wide area network, the audio information buffer 31 will receive and buffer audio
10 and playback timing information from the network communications manager 40. It will be
11 appreciated that, if the zone player 11(n) is not a member of the synchrony group, the zone player
12 11(n) will not play this buffered audio and playback timing information.

13 On yet another hand, if the zone player 11(n) is operating as the master device 21 or a slave
14 device 22(g) in a synchrony group, and if the zone player 11(n) is not also the audio information
15 channel device 23 providing audio and playback timing information for the synchrony group 20, the
16 audio information buffer 31 will receive and buffer audio and playback timing information from the
17 network communications manager 40.

18 The audio information source interface 30 connects to the audio information source(s)
19 14(n)(s) associated with the zone player 11(n). While the zone player 11(n) is operating as audio
20 information channel device 23 for a synchrony group 20, and if the audio information is to be
21 provided by a source 14(n)(s) connected to the zone player 11(n), the audio information source
22 interface 30 will selectively receive audio information from one of the audio information source(s)
23 14(n)(s) to which the zone player is connected and store the audio information in the audio
24 information buffer 21. If the audio information from the selected audio information source 14(n)(s)
25 is in analog form, the audio information source interface 30 will convert it to digital form. The
26 selection of the audio information source 14(n)(s) from which the audio information source interface
27 30 receives audio information is under control of the control module 42, which, in turn, receives
28 control information from the user interface module through the user interface module interface 44.

1 The audio information source interface 30 adds playback timing information to the digital audio
2 information and buffers the combined audio and playback timing information in the audio
3 information buffer 21.

4 More specifically, as noted above, the audio information source interface 30 receives audio
5 information from an audio information source 14(n)(s), converts it to digital form if necessary, and
6 buffers it along with playback timing information in the audio information buffer 21. In addition,
7 the audio information source interface 30 will also provide formatting and scheduling information
8 for the digital audio information, whether as received from the selected audio information source
9 14(n)(s) or as converted from an analog audio information source. As will be made clear below, the
10 formatting and scheduling information will control not only playback by the zone player 11(n) itself,
11 but will also enable other zone players 11(n'), 11(n''),... that may be in a synchrony group for which
12 the zone player 11(n) is the master device, to play the audio program associated with the audio
13 information in synchrony with the zone player 11(n).

14 In one particular embodiment, the audio information source interface 30 divides the audio
15 information associated with an audio work into a series of frames, with each frame comprising
16 digital audio information for a predetermined period of time. As used herein, an audio track may
17 comprise any unit of audio information that is to be played without interruption. On the other hand,
18 an audio program may comprise a series of one or more audio tracks that are to be played in
19 succession. It will be appreciated that the tracks comprising the audio program may also be played
20 without interruption, or alternatively playback between tracks may be interrupted by a selected time
21 interval. FIG. 4 schematically depicts an illustrative framing strategy used in connection with one
22 embodiment of the invention for a digital audio stream comprising an audio work. More
23 specifically, FIG. 4 depicts a framed digital audio stream 50 comprising a sequence of frames 51(1)
24 through 51(F) (generally identified by reference numeral 51(f)). Each frame 51(f), in turn, comprises
25 a series of audio samples 52(f)(1) through 52(f)(S) (generally identified by reference numeral
26 52(f)(s)) of the audio track. Preferably all of the frames will have the same number "S" of audio
27 samples, although it will be appreciated from the following that that is primarily for convenience.
28 On the other hand, it will be appreciated that, the number of audio samples may differ from "S"; this

1 may particularly be the case if the frame 51(f) contains the last audio samples for the digital audio
2 stream for a particular audio work. In that case, the last frame 51(F) will preferably contain samples
3 52(F)(1) through 52(F)(x), where "x" is less than "S." Generally, it is desirable that the number of
4 samples be consistent among all frames 51(f), and in that case padding, which will not be played,
5 can be added to the last frame 51(F).

6 Associated with each frame 51(f) is a header 55(f) that includes a number of fields for storing
7 other information that is useful in controlling playback of the audio samples in the respective frame
8 51(f). In particular, the header 55(f) associated with a frame 51(f) includes a frame sequence number
9 field 56, an encoding type field 57, a sampling rate information field 58, a time stamp field 60, an
10 end of track flag 61, and a length flag field 62. The header 55(f) may also include fields (not shown)
11 for storing other information that is useful in controlling playback. Generally, the frame sequence
12 number field 56 receives a sequence number "f" that identifies the relative position of the frame 51(f)
13 in the sequence of frames 51(1)..51(f)...51(F) containing the digital audio stream 50. The encoding
14 type field 57 receives a value that identifies the type of encoding and/or compression that has been
15 used in generating the digital audio stream. Conventional encoding or compression schemes include,
16 for example, the well-known MP3 and WAV encoding and/or compression schemes, although it will
17 be appreciated that other schemes may be provided for as well. The sampling rate information field
18 58 receives sampling rate information that indicates the sampling rate for the audio samples 52(f)(s).
19 As will be apparent to those skilled in the art, the sampling rate determines the rate at which the zone
20 player 11(n) is to play the audio samples 52(f)(s) in the frame, and, as will be described below,
21 determines the period of the digital to analog converter clock 34.

22 The condition of the end of work flag 61 indicates whether the frame 51(f) contains the last
23 digital audio samples for the audio track associated with the framed digital audio work 50. If the
24 frame 51(f) does not contain the audio samples that are associated with the end of the digital audio
25 stream 50 for a respective audio work, the end of work flag will be clear. On the other hand, if the
26 frame 51(f) does contain the audio samples that are associated with the end of the digital audio
27 stream 50 for a respective audio work, the end of work flag 61 will be set. In addition, since the
28 number of valid audio samples 52(F)(s) in the frame 51(F), that is, the samples that are not padding,

1 may be less than "S," the default number of audio samples in a frame 51(f), the length flag field 62
2 will contain a value that identifies the number of audio samples in 52(F)(s) in the last frame 51(F)
3 of the audio work 50. If, as noted above, the frames have a consistent number "S" of samples, the
4 samples 52(F)(x+1) through 52(F)(S) will contain padding, which will not be played.

5 The time stamp field 60 stores a time stamp that identifies the time at which the zone player
6 11(n) is to play the respective frame. More specifically, for each frame of a framed digital audio
7 stream 50 that is buffered in the audio information buffer 21, the audio information source interface
8 30, using timing information from the digital to analog converter clock 34, will determine a time at
9 which the zone player 11(n) is to play the respective frame, and stores a time stamp identifying the
10 playback time in the time stamp field 60. The time stamp associated with each frame will later be
11 used by the playback scheduler 32 to determine when the portion of the digital audio stream stored
12 in the frame is to be coupled to the digital to analog converter 33 to initiate play back. It will be
13 appreciated that the time stamps that are associated with frames in sequential frames 51(1),
14 51(2),...,51(F), will be such that they will be played back in order, and without an interruption
15 between the sequential frames comprising the digital audio stream 50. It will further be appreciated
16 that, after a time stamp has been determined for the first frame, stored in frame 51(1), of a digital
17 audio stream 50, the audio information source interface 30 can determine time stamps for the
18 subsequent frame 51(2), 51(3),...,51(F) in relation to the number of samples "S" in the respective
19 frames and the sample rate. The time stamps will also preferably be such that frames will be played
20 back after some slight time delay after they have been buffered in the audio information buffer 21;
21 the purpose for the time delay will be made clear below.

22 Returning to FIG. 3, in addition to dividing the digital audio information into frames, the
23 audio information source interface 30 also aggregates and/or divides the frames 51(f) as necessary
24 into packets, each of which will be of a length that would fit into a message for transmission over
25 the network, and associates each packet with a packet sequence number. For example, if a packet
26 will accommodate multiple frames 51(f), 51(f+1),... 51(f+y-1), it will aggregate them into a packet
27 and associate them with a packet number, for example p(x). If the entire frames 51(f) and 51(f+y-1)
28 was accommodated in packet p(x), where "x" is the sequence number, which will occur if the size

1 of a packet is an exact multiple of the frame size, the next packet, $p(x+1)$ will begin with frame
2 $51(f+y)$ and will include frames $51(f+y), \dots, 51(f+2y-1)$. Subsequent packets $p(x+2), \dots$ will be
3 formed in a similar manner. On the other hand, if the packet length will not accommodate an exact
4 multiple of the frame size, the last frame in the packet will be continued at the beginning of the next
5 packet.

6 If the audio information source interface 30 is aware of track boundaries, which may be the
7 case if the tracks are divided into files, the packets will reflect the track boundaries, that is, the
8 packets will not contain frames from two tracks. Thus, if the last frames associated with a track are
9 insufficient to fill a packet, the packet will contain padding from the last frame associated with the
10 track to the end of the packet, and the next packet will begin with the first frames associated with
11 the next track.

12 In one embodiment, the audio information source interface 30 stores the packets in the audio
13 information buffer 31 in a ring buffer. As is conventional, a ring buffer includes a series of storage
14 locations in the buffer. Each entry will be sufficient to store one packet. Four pointers are used in
15 connection with the ring buffer, a first pointer pointing to the beginning of the ring buffer, a second
16 pointer pointing to the end of the ring buffer, a third "write" pointer pointing to the entry into which
17 a packet will be written and a fourth "read" pointer pointing to the entry from which packet will be
18 read for use in playback. When a packet is read from the ring buffer for playback, it will be read
19 from the entry pointed to by the read pointer. After the packet has been read, the read pointer will
20 be advanced. If the read pointer points beyond the end of the ring buffer, as indicated by the end
21 pointer, it will be reset to point to the entry pointed to by the beginning pointer, and the operations
22 can be repeated.

23 On the other hand, when the audio information source interface 30 stores a packet in the ring
24 buffer, first determine whether the entry pointed to by the write pointer points to the same entry as
25 the entry pointed to by the read pointer. If the write pointer points to the same entry as the entry
26 pointed to by the read pointer, the entry contains at least a portion of a packet that has not yet been
27 read for playback, and the audio information source interface 30 will delay storage of the packet until

1 the entire packet has been read and the read pointer advanced. After the read pointer has been
2 advanced, the audio information source interface 30 can store the packet in the entry pointed to by
3 the write pointer. After the packet has been stored, the audio information source interface 30 will
4 advance the write pointer. If the write pointer points beyond the end of the ring buffer, as indicated
5 by the end pointer, it will be reset to point to the entry pointed to by the beginning pointer, and the
6 operations can be repeated.

7 As noted above, the zone player 11(n) can operate both as an audio information channel
8 device 23 and a member of the synchrony group 20 of which it is a member. In that case, the audio
9 information buffer 31 can contain one ring buffer. On the other hand, the zone player 11(n) can
10 operate as an audio information channel device 23 for one synchrony group 20(1) (FIG. 2A) and a
11 member of another synchrony group 20(2). In that case, the audio information buffer 31 would
12 maintain two ring buffers, one for the audio and timing information associated with synchrony group
13 20(1), and the other for the audio and timing information associated with synchrony group 20(2).
14 It will be appreciated that, in the latter case, the zone player 11(n) will only use the audio and timing
15 information that is associated with synchrony group 20(2) for playback.

16 The playback scheduler 32 schedules playback of the audio information that is buffered in
17 the audio information buffer 31 that is to be played by the zone player 11(n). Accordingly, under
18 control of the playback scheduler 32, the digital audio information that is buffered in the audio
19 information buffer 21 that is to be played by the zone player 11(n) is transferred to the digital to
20 analog converter 33 for playback. As noted above, if the zone player 11(n) is operating as an audio
21 information channel device 23 for a synchrony group 20 for which it is not a member, the playback
22 scheduler 32 will not schedule the digital audio information that is to be played by that synchrony
23 group 20 for playback. The playback scheduler 32 only schedules the digital audio information, if
24 any, that is buffered in the audio information buffer 31 that is associated with a synchrony group for
25 which the zone player 11(n) is a member, whether as master device 21 or a slave device 22(g).

26 Essentially, the playback scheduler 32 makes use of the read pointer associated with the
27 circular buffer that contains the audio and playback timing information that is to be played by the

1 zone player 11(n). The playback scheduler 32 retrieves the packet information from the entry of the
2 ring buffer pointed to by the read pointer, and then advances the ring pointer as described above. The
3 playback scheduler 32 determines the boundaries of the frames in the packet and uses the time
4 stamps in the time stamp fields 60 associated with the respective frame 51(f), along with timing
5 information provided by the zone player 11(n)'s digital to analog converter clock 34, to determine
6 when the respective frame is to be transferred to the digital to analog converter 33. Generally, when
7 the time stamp associated with a buffered digital audio information frame corresponds to the current
8 time as indicated by the digital to analog converter clock 34, the playback scheduler 32 will enable
9 the respective frame to be transferred to the digital to analog converter 33.

10 The digital to analog converter 33, also under control of the digital to analog converter clock
11 34, converts the buffered digital audio information to analog form, and provides the analog audio
12 information to the audio amplifier 35 for amplification. The amplified analog information, in turn,
13 is provided to the audio reproduction devices 15(n)(r) through the audio reproduction device
14 interface 36. The audio reproduction devices 15(n)(r) transform the analog audio information signal
15 to sound thereby to provide the audio program to a listener. The amount by which the audio
16 amplifier 35 amplifies the analog signal is controlled by the control module 42, in response to
17 volume control information provided by the user through the user interface module 13.

18 The network communications manager 40 controls network communications over the
19 network 12, and the network interface 41 transmits and receives message packets over the network
20 12. The network communications manager 40 generates and receives messages to facilitate the
21 transfer of the various types of information described above in connection with FIG. 2, including the
22 channel device control information, slave device control information, audio and playback timing
23 information and the audio information channel device's clock timing information. In connection
24 with the channel device control information and the slave device control information, the network
25 communications manager 40 will generate messages for transfer over the network 12 in response to
26 control information from the control module 42. Similarly, when the network communications
27 manager 40 receives messages containing channel device control information and slave device

1 control information, the network communications manager will provide the information to the
2 control module 42 for processing.

3 With regards to the audio information channel device's clock timing information, as noted
4 above, the master device 21 and slave devices 22(g) of the synchrony group 20 obtain the clock
5 timing information from the audio information channel device 23 using the well-known SNTP. If
6 the zone player 11(n) is operating as the audio information channel device 23 for a synchrony group,
7 during the SNTP operation, it will provide its current time, particularly a current time as indicated
8 by its digital to analog converter clock 34. On the other hand, if the zone player 11(n) is operating
9 as the master device 21 or slave device 22(g) of a synchrony group 20, it will receive the clock
10 timing information from the audio information channel device 23. After the respective device 21,
11 22(g) has obtained the audio information channel device's clock timing information, it will generate
12 a differential time value ΔT representing the difference between the time T indicated by its digital
13 to analog converter clock 34 and the current time information from the audio information channel
14 device 23. The differential time value will be used to update the time stamps for the frames of the
15 digital audio stream 50 (FIG. 4) that are received from the audio information channel device.

16 With regards to the audio and playback timing information, operations performed by the
17 network communications manager 40 will depend on whether

18 (i) the audio and playback timing information has been buffered in the audio information
19 buffer 31 for transmission, as audio information channel device 23, over the network 12 to the
20 master device 21 and/or slave devices 22(g) of a synchrony group, or

21 (ii) the audio and playback timing information has been received from the network 12 to be
22 played by the zone player 11(n) as either the master device 21 for a synchrony group or a slave
23 device in a synchrony group.

24 It will be appreciated that the network communications manager 40 may be engaged in both (i) and
25 (ii) contemporaneously, since the zone player 11(n) may operate both as the audio information
26 channel device 23(1) for a synchrony group 20(1) (reference FIG. 2A) of which it is not a member,

1 and a member of another synchrony group 20(2) for which another zone player 11(n') is the audio
2 information channel device 23(2). With reference to item (i) above, after a packet that is to be
3 transmitted has been buffered in the respective ring buffer, the network communications manager
4 40 retrieves the packet, packages it into a message and enables the network interface 41 to transmit
5 the message over the network 12. If the control module 42 receives control information from the
6 user interface module 13 (if the master device 21 is also the audio information channel device 23 for
7 the synchrony group 20) or from the master device (if the master device 21 is not the audio
8 information channel device 23 for the synchrony group 20) that would require the transmission of
9 a "resynchronize" command as described above, the control module 42 of the audio information
10 channel device 23 enables the network communications manager 40 to insert the command into a
11 message containing the audio and playback timing information. Details of the operations performed
12 in connection with the "resynchronize" command will be described below. As noted above, the
13 "resynchronize" command is used if the user enables a synchrony group to terminate the playback
14 of a track that is currently being played, or cancel playback of a track whose playback has not begun.

15 On the other hand, with reference to item (ii) above, if network interface 41 receives a
16 message containing a packet containing frames of audio and playback timing information that the
17 zone player 11(n) is to play either as a master device 21 or a slave device for a synchrony group 20,
18 the network interface 41 provides the audio and playback timing information to the network
19 communications manager 40. The network communications manager 40 will determine whether the
20 packet contains a resynchronize command and, if so, notify the control module 42, which will enable
21 operations to be performed as described below. In any case, the network communications manager
22 40 will normally buffer the various frames comprising the audio and playback timing information
23 in the audio information buffer 31, and in that operation will generally operate as described above
24 in connection with the audio information source interface 30. Before buffering them, however, the
25 network communications manager 40 will update their time stamps using the time differential value
26 described above. It will be appreciated that the network communications manager 40 will perform
27 similar operations whether the messages that contain the packets were multi-cast messages or unicast
28 messages as described above

1 The updating of the time stamps by the master device 21 and the slave devices 22(g) in the
 2 synchrony group 20 will ensure that they all play the audio information synchronously. In particular,
 3 after the network communications manager 40 has received a frame 51(f) from the network interface
 4 41, it will also obtain, from the digital to analog converter clock 34, the zone player 11(n)'s current
 5 time as indicated by its digital to analog converter clock 34. The network communications manager
 6 40 will determine a time differential value that is the difference between the slave device's current
 7 clock time, as indicated by its digital to analog converter 34, and the audio information channel
 8 device's time as indicated by the audio information channel device's clock timing information.
 9 Accordingly, if the master or slave device's current time has a value T_S and the audio information
 10 channel device's current time, as indicated by the clock timing information, has a value T_C , the time
 11 differential value $\Delta T = T_S - T_C$. If the current time of the master or slave device in the synchrony group
 12 20, as indicated by its digital to analog converter clock 34, is ahead of the audio information channel
 13 device's clock time as indicated by the clock timing information received during the SNTP operation,
 14 the time differential value will have a positive value. On the other hand, if the master or slave
 15 device's current time is behind the audio information channel device's clock time, the time
 16 differential value ΔT will have a negative value. If the zone player 11(n) obtains clock timing
 17 information from the audio information channel device 23 periodically while it is a member of the
 18 synchrony group 20, the network communications manager 40 can generate an updated value for the
 19 time differential value ΔT when it receives the clock timing information from the audio information
 20 channel device 23, and will subsequently use the updated time differential value.

21 The network communications manager 40 uses the time differential value ΔT that it generates
 22 from the audio information channel device timing information and zone player 11(n)'s current time
 23 to update the time stamps that will be associated with the digital audio information frames that the
 24 zone player 11(n) receives from the audio information channel device. For each digital audio
 25 information frame that is received from the audio information channel device, instead of storing the
 26 time stamp that is associated with the frame as received in the message in the audio information
 27 buffer 21, the network communications manager 40 will store the updated time stamp with the
 28 digital audio information frame. The updated time stamp is generated in a manner so that, when the

1 zone player 11(n), as a member of the synchrony group plays back the digital audio information
2 frame, it will do so in synchrony with other devices in the synchrony group.

3 More specifically, after the zone player 11(n)'s network interface 41 receives a message
4 containing a packet that, in turn, contains one or more frames 51(f), it will provide the packet to the
5 network communications manager 40. For each frame 51(f) in the packet that the network
6 communications manager 40 receives from the network interface 41, the network communications
7 manager 40 will add the time differential value ΔT to the frame's time stamp, to generate the updated
8 time stamp for the frame 51(f), and store the frame 51(f), along with the header 55(f) with the
9 updated time stamp in the audio information buffer 31. Thus, for example, if a frame's time stamp
10 has a time value T_F , the network communications manager 40 will generate an updated time stamp
11 T_F^U having a time value $T_F^U = T_F + \Delta T$. Since time value T_F^U according to the slave device's digital
12 to analog converter clock 34 is simultaneous to the time value T_F according to the audio information
13 channel device's digital to analog converter clock 34, the zone player 11(n) device will play the
14 digital audio information frame at the time determined by the audio information channel device 23.
15 Since all of the members of the synchrony group 20 will perform the same operations, generating
16 the updated time stamps T_F^U for the various frames 51(f) in relation to their respective differential
17 time values, all of the zone players 11(n) in the synchrony group 20 will play them synchronously.
18 The network communications manager 40 will generate updated time stamps T_F^U for all of the time
19 stamps 60 in the packet, and then store the packet in the audio information buffer 31.

20 It will be appreciated that, before storing a packet in the audio information buffer 21, the
21 network communications manager 40 can compare the updated time stamps T_F^U associated with the
22 frames in the packet to the slave device's current time as indicated by its digital to analog converter
23 clock 34. If the network communications manager 40 determines the time indicated by the updated
24 time stamps of frames 51(f) in the packet are earlier than the zone player's current time, it can discard
25 the packet instead of storing it in the audio information buffer 21, since the zone player 11(n) will
26 not play them. That is, if the updated time stamp has a time value T_F^U that identifies a time that is
27 earlier than the zone player's current time T_S as indicated by the zone player's digital to analog
28 converter clock 34, the network communications manager 40 can discard the packet.

1 If the zone player 11(n) is operating as the master device 21 of a synchrony group 20, when
2 the user, through the user interface module 13, notifies the zone player 11(n) that another zone player
3 11(n') is to join the synchrony group 20 as a slave device 22(g), the control module 42 of the zone
4 player 11(n) enables the network communications manager 40 to engage in an exchange of
5 messages, described above in connection with FIG. 2, to enable the other zone player 11(n') to join
6 the synchrony group 20 as a slave device. As noted above, during the message exchange, the
7 messages generated by the network communications manager 40 of the zone player 11(n) will
8 provide the network communications manager of the zone player 11(n') that is to join the synchrony
9 group 20 with information such as the multi-cast address being used by the audio information
10 channel device 23 that is providing the audio program to the synchrony group 20, as well as a unicast
11 network address for the audio information channel device 23. After receiving that information, the
12 network communications manager and network interface of the zone player 11(n') that is to join the
13 synchrony group 20 can begin receiving the multi-cast messages containing the audio program for
14 the synchrony group, engage in SNTP transactions with the audio information channel device 23 to
15 obtain the latter's current time, and also enable the audio information channel device 23 to send the
16 zone player 11(n') frames 51(f) that it had previously broadcast using the unicast message
17 transmission methodology as described above.

18 On the other hand, if the network communications manager 40 and network interface 41 of
19 the zone player 11(n) receive a message over the network 12 indicating that it is to become a slave
20 device 22(g) of a synchrony group for which another zone player 11(n') is the master device, the
21 network communications manager 40 for zone player 11(n) will provide a notification to the control
22 module 42 of zone player 11(n). Thereafter, the control module 42 of zone player 11(n) can enable
23 the network communications manager 40 of zone player 11(n) to perform the operations described
24 above to enable it to join the synchrony group 20.

25 As noted above, the user, using user interface module 13, can enable the synchrony group
26 to terminate playback of a track of an audio program that is currently being played. After playback
27 of a track that is currently being played has been terminated, playback will continue in a
28 conventional manner with the next track that has been buffered in the audio information buffer 31.

1 It will be appreciated that that could be the next track that is on the original play list, or a previous
2 track. In addition, the user can enable the synchrony group 20 to cancel playback of a track that it
3 has not yet begun to play, but for which buffering of packets has begun in the synchrony group 20.
4 Both of these operations make use of the "resynchronize" command that the master device 21 of the
5 synchrony group 20 can enable the audio information channel device 23 to include in the multi-cast
6 message stream that it transmits to the synchrony group 20. Generally, in response to receipt of the
7 resynchronize command, the members of the synchrony group 20 flush the ring buffer containing
8 the packets that they are to play in the future. In addition, if the members of the synchrony group
9 provide separate buffers for their respective digital to analog converters 33, the members will also
10 flush those buffers as well. After the audio information channel device transmits a packet containing
11 the resynchronize command:

12 (i) in the case of the use of the resynchronize command to terminate playing of a track that
13 is currently being played, the audio information channel device 23 will begin multi-casting packets
14 for the next track, to begin play immediately, and will continue through the play list in the manner
15 described above; and

16 (ii) in the case of the use of the resynchronize command to cancel play of a track for which
17 buffering has begun, but which is to be played in the future, the audio information channel device
18 23 will begin multi-casting packets for the track after the track that has been cancelled, to be played
19 beginning at the time the cancelled track was to begin play, and will also continue through the play
20 list in the manner described above.

21 It will be appreciated that,

22 (a) in the first case (item (i) directly above), the resynchronize command can enable the read
23 pointer to be set to point to the entry in the circular buffer into which the first packet for the next
24 track will be written, which will correspond to the entry to which the write pointer points, but

(b) in the second case (item (ii) directly above), the resynchronize command can enable the write pointer for the circular buffer to be set to point to the entry that contains the first packet for the track whose play is being cancelled.

It will further be appreciated that, if a track is cancelled for which buffering has not begun, the resynchronize command will generally not be necessary, since the audio information channel device 23 for the synchrony group 20 need only delete that track from the play list.

Operations performed in connection with use of the resynchronize command to cancel playback of a track that is currently being played will be described in connection with Packet Sequence A below, and operations performed in connection with the use of the resynchronize command to cancel playback of a track that is has not yet begun to play, but for which buffering of packets has begun, will be described in connection with Packet Sequence B below.

Packet Sequence A

(A1.0) [packet 57]

(A1.1) [continuation of frame 99]

(A1.2) [frame 100, time = 0:00:01, type = mp3 audio]

(A1.3) [frame 101, time = 0:00:02, type = mp3 audio]

(A1.4) [frame 102, time = 0:00:03, type = mp3 audio]

(A2.0) [packet 58]

(A2.1) [continuation of frame 102]

(A2.2) [frame 103, time = 0:00:04, type = mp3 audio]

(A2.3) [frame 104, time = 0:00:05, type = mp3 audio]

(A2.4) [frame 105, time = 0:00:06, type = mp3 audio]

(A3.0) [packet 59]

(A3.1) [continuation of frame 105]

(A3.2) [frame 106, time = 0:00:07, type = mp3 audio]

- 1 (A3.3) [frame 107, time = 0:00:08, type = mp3 audio]
- 2 (A3.4) [frame 108, time = 0:00:09, type = mp3 audio]
- 3
- 4 (A4.0) [packet 60]
- 5 (A4.1) [continuation of frame 108]
- 6 (A4.2) [frame 109, time = 0:00:10, type = mp3 audio]
- 7 (A4.3) [Resynchronize command]
- 8 (A4.4) [Padding, if necessary]
- 9
- 10 (A5.0) [packet 61]
- 11 (A5.1) [frame 1, time = 0:00:07, type = mp3 audio]
- 12 (A5.2) [frame 2, time = 0:00:08, type = mp3 audio]
- 13 (A5.3) [frame 3, time = 0:00:09, type = mp3 audio]
- 14 (A5.4) [frame 4, time = 0.00.10, type = mp3 audio]
- 15
- 16 (A6.0) [packet 62]
- 17 (A6.1) [continuation of frame 4]
- 18 (A6.2) [frame 5, time = 0:00:11, type = mp3 audio]
- 19 (A6.3) [frame 6, time = 0:00:12, type = mp3 audio]
- 20 (A6.4) [frame 7, time = 0:00:13, type = mp3 audio]

21 Packet Sequence A comprises a sequence of six packets, identified by packet 57 through
 22 packet 62, that the audio information channel device 23 multi-casts in respective messages to the
 23 members of a synchrony group 20. It will be appreciated that the series of messages that the audio
 24 information channel device 23 may multi-cast to the synchrony group 20 may include a messages
 25 prior to the packet 57, and may also include messages after packet 62. Each packet comprises a
 26 packet header, which is symbolized by lines (A1.0), (A2.0),...(A6.0) in Packet Sequence A, and will
 27 generally also include information associated with at least a portion of a frame. In the packets
 28 represented in Packet Sequence A, each packet includes information associated with a plurality of

1 frames. Depending on the lengths of the packets, each packet may contain information associated
2 with a portion of a frame, an entire frame, or multiple frames. In the illustration represented by
3 Packet Sequence A, it is assumed that each packet may contain information associated with multiple
4 frames. In addition, it is assumed that a packet does not necessarily contain information associated
5 with an integral number of frames; in that case, a packet may contain information associated with
6 a portion of a frame, and the next packet will contain the information associated with the rest of a
7 frame.

8 The frames and associated header playback timing information contained in the various
9 packets are symbolized by lines (A1.1), (A1.2),..., (A1.4), (A2.1),..., (A6.4) of Packet Sequence A.
10 Thus, for example, line (A1.2) of packet 57 represents the one-hundredth frame, that is, frame
11 51(100) (reference FIG. 4), of the track whose audio information is being transmitted in the sequence
12 of packets that includes packet 57. The frame 51(100) is to be played at an illustrative time,
13 according to the audio information channel device's digital to analog converter clock, of
14 "time = 0:00:01," and the frame is encoded and/or compressed using the well-known MP3 encoding
15 and compression methodology. In that case, the legend "time = 0:00:01" represents the time stamp
16 that would be included in field 60 (FIG. 4) of the header associated with the frame 50(100) as multi-
17 cast by the audio information channel device for the synchrony group. It will be appreciated that the
18 playback time and encoding/compression methodology will be referred in the header 55(100) that
19 is associated with the frame 51(100). It will also be appreciated that the header may also contain
20 additional information as described above.

21 Similarly, line (A1.3) of packet 57 represents the one-hundred and first frame, that is, frame
22 51(101), of the track whose audio information is being transmitted in the sequence of packets that
23 includes packet 57. The frame 51(101) is to be played at an illustrative time, according to the audio
24 information channel device's digital to analog converter clock, of "0:00:02," and the frame is also
25 encoded and/or compressed using the MP3 encoding and compression methodology. Line (A1.4)
26 of packet 57 represents similar information, although it will be appreciated that, depending on the
27 length of packet 57, the line may not represent the information for an entire frame 51(102) and/or
28 its associated header. If the length of packet 57 is not sufficient to accommodate the information for

1 the entire frame 51(102) and/or associated header, the information will continue in packet 58, as
2 represented by line (A2.1) in Packet Sequence A. Similarly, if the length of packet 56 was not
3 sufficient to contain the information for an entire frame 51(99) preceding frame 51(100), packet 57
4 (lines (A1.0) through 1.4) may contain any information from frame 51(99) that packet 56 was unable
5 to accommodate.

6 As noted above, when the master device 21 or a slave device 22(g) in the synchrony group
7 20 receives the packet 57, its respective network communications manager 40 will update the time
8 stamps associated with the various frames 51(f) as described above before buffering the respective
9 frames in the respective audio information buffer 31.

10 Packets 58 and 59 contain information that is organized along the lines described above in
11 connection with packet 57.

12 Packet 60 also contains, as represented by lines (A4.1) and (A4.2), information that is
13 organized along the lines of the information represented by lines (Ax.1) and (Ax.2) ("x" equals an
14 integer) described above in connection with packets 57 through 59. On the other hand, packet 60
15 contains a resynchronize command, as represented by line (A4.3). Packet 60 also may contain
16 padding, as represented by line 4.4, following the resynchronize command. As noted above, the
17 master device 21 of a synchrony group 20 will enable the audio information channel device 23 that
18 is providing audio information to the synchrony group 20 to multi-cast a message containing the
19 resynchronize command when it receives notification from the user interface module 13 that the user
20 wishes to cancel playback of a track that is currently being played. In the example depicted in
21 Packet Sequence A, as will be described below, the audio information channel device 23 receives
22 notification from the master device 21 that the user wishes to cancel playback of a track at a time
23 corresponding to "time = 0:00:07" according to its digital to analog converter clock 34, and, in line
24 (A4.3) of packet 60 it will provide the resynchronize command, followed by padding, if necessary.

25 As will be apparent from examining lines (A3.1) through (A3.4) of packet 59 and lines
26 (A4.1) and (A4.2) of packet 60, although the audio information channel device 23 has received the
27 notification from the synchrony group's master device 21 to multi-cast the resynchronize command

1 at a time corresponding to "time = 0:00:07" according to the clock time indicated by its digital to
2 analog converter clock 34, it (that is, the audio information channel device 23) has already multi-cast
3 messages containing frames that are to be played at that time and subsequently. That is, the audio
4 information channel device 23 has, multi-cast in packet 59, frames 51(106) through 51(108) that
5 contain time stamps "time = 0:00:07," "time = 0:00:08" and "time = 0:00:09," respectively, and, in
6 packet 60, in addition to the continuation of frame 51(108), frame 51(109) that contains time stamp
7 "time = 0:00:10." (It will be appreciated that the times indicated by the illustrative time stamps are
8 for illustration purposes only, and that in an actual embodiment the time stamps may have different
9 values and differentials.)

10 As noted above, the audio information channel device 23 multi-casts a message containing
11 a packet that, in turn, contains the resynchronize command when it receives the notification from
12 the master device 21 to do so. In the example depicted in Packet Sequence A, the packet will be
13 multi-cast when the audio information channel device's digital to analog converter clock time
14 corresponds to "0:00:07." Subsequently, two things happen. In one, aspect, when the master device
15 21 and slave devices 22(g) receive the packet that contains the resynchronize command, they will
16 stop playing the audio program that they are playing.

17 In addition, the audio information channel device 23 will begin transmitting frames
18 containing audio information for the next track, including therewith time stamps immediately
19 following the digital to analog converter clock time at which the packet including the resynchronize
20 command was transmitted. Accordingly, and with further reference to Packet Sequence A, the audio
21 information channel device 23 will multi-cast a message containing packet 61. As indicated above,
22 packet 61 contains, as represented in lines (A5.1) through (A5.3), frames 51(1) through 51(3), which
23 are the first three frames of the next track of the audio program that is to be played. It is also
24 compressed and encoded using the MP3 encoding and compression scheme, and it is accompanied
25 by time stamps "time = 0:00:07," "time = 0:00:08" and "time = 0:00:10." As noted above, the time
26 stamp "time = 0:00:07" corresponds to the clock time at which the audio information channel device
27 23 multi-casts the resynchronize command, and, when the master device 21 and slave devices 22(g)
28 receive these frames, they would be expected to begin playing them very shortly, if not immediately

after the audio information channel device 23 multi-casts the message containing the packet that, in turn, contains the resynchronize command. Packet 61 also includes at least a portion of the next frame, that is, frame 51(4), for that track. In addition, Packet Sequence A depicted above further includes a subsequent packet, namely, packet 62, that contains any necessary continuation of frame 51(4), as well as three subsequent frames. If any additional packets are required for the track, as well as for subsequent tracks, they can be multi-cast in a similar manner.

As further noted above, the resynchronize command can also be used to cancel playing of one or more tracks for which playback has begun. This will be illustrated in connection with Packet Sequence B:

Packet Sequence B

(B1.0) [packet 157]

(B1.1) [continuation of frame 99]

(B1.2) [frame 100, time = 0:00:01, type = mp3 audio]

(B1.3) [frame 101, time = 0:00:02, type = mp3 audio]

(B1.4) [frame 102, time = 0:00:03, type = mp3 audio]

(B2.0) [packet 158]

(B2.1) [continuation of frame 102]

(B2.2) [frame 103, time = 0:00:04, type = mp3 audio]

(B2.3) [frame 104, time = 0:00:05, type = mp3 audio]

(B2.4) [frame 105, time = 0:00:06, type = mp3 audio]

(B3.0) [packet 159]

(B3.1) [continuation of frame 105]

(B3.2) [frame 106, time = 0:00:07, type = mp3 audio]

(B3.3) [track boundary notification]

(B3.4) [Padding, if necessary]

1
2 (B4.0) [packet 160]
3 (B4.1) [frame 1, time = 0:00:08, type = mp3 audio]
4 (B4.2) [frame 2, time = 0:00:09, type = mp3 audio]
5 (B4.3) [frame 3, time = 0:00:10, type = mp3 audio]
6
7 (B5.0) [packet 161]
8 (B5.1) [continuation of frame 3]
9 (B5.2) [frame 4, time = 0:00:11, type = mp3 audio]
10 (B5.3) [Resynchronize, after packet 159]
11 (B5.4) [Padding, if necessary]
12
13 (B6.0) [packet 162]
14 (B6.1) [frame 1, time = 0:00:08, type = mp3 audio]
15 (B6.2) [frame 2, time = 0:00:09, type = mp3 audio]
16 (B6.3) [frame 3, time = 0:00:10, type = mp3 audio]
17 (B6.4) [frame 4, time = 0:00:11, type = mp3 audio]
18
19 (B7.0) [packet 163]
20 (B7.1) [continuation of frame 4]
21 (B7.2) [frame 5, time = 0:00:12, type = mp3 audio]
22 (B7.3) [frame 6, time = 0:00:13, type = mp3 audio]
23 (B7.4) [frame 7, time = 0:00:14, type = mp3 audio]

24 Packet Sequence B comprises a series of seven packets, identified by packet 157 through 163
25 , that the audio information channel device 23 multi-casts to the members of a synchrony group 20.
26 As with Packet Sequence A, it will be appreciated that the series of packets that the audio
27 information channel device 23 may multi-cast to the synchrony group 20 may include packets prior
28 to the packet 157, and may also include packets after packet 162. Each packet comprises a packet

1 header, which is symbolized by lines (B1.0), (B2.0),...(B7.0) in Packet Sequence B. As in Packet
2 Sequence A, each packet will also generally include information associated with at least a portion
3 of a frame 51(f) along with its associated frame 55(f). As in the packets represented in Packet
4 Sequence A, each packet includes information associated with a plurality of frames. Depending on
5 the lengths of the packets, each packet may contain information associated with a portion of a frame,
6 an entire frame, or multiple frames. Further, as with Packet Sequence A, it is assumed that each
7 packet may contain information associated with multiple frames. In addition, it is assumed that a
8 packet does not necessarily contain information associated with an integral number of frames; in that
9 case, a packet may contain information associated with a portion of a frame, and the next packet will
10 contain the information associated with the rest of a frame.

11 The structures of the packets represented by Packet Sequence B are similar to those described
12 above in connection with Packet Sequence A, and will not be repeated here. Generally, Packet
13 Sequence B illustratively contains a sequence of packets that represent at least portions of three
14 tracks that may have been selected from, for example, a play list. In particular, packets 157 through
15 159 represent frames from a portion of one track, packets 160 and 161 represent frames from a
16 second track and packets 162 and 163 represent frames from a third track. The play list indicated
17 that the first, second and third tracks were to be played in that order. With particular reference to
18 Packet Sequence B, it should be noted that line (B3.3) indicates that packet 159 includes an
19 indication that that packet contains the last frame for the track, and line (B3.4) provides for padding
20 to the end of the packet. The first frame of the next track begins in packet 160.

21 In connection with the use of the resynchronize command to cancel playback of a track, at
22 least a portion of which the audio information channel device 23 has multi-cast to the members of
23 the synchrony group, packet 161, in line (B5.3) represents a resynchronize command that indicates
24 that resynchronization is to occur after packet 159, that is, immediately after the packet that contains
25 the last frame of the first of the three tracks represented by the packets in Packet Sequence B. It
26 should be noted that the resynchronize command is in the packet 161, while the resynchronization
27 is to occur at packet 160, that is, the synchrony group is to not play the track starting with packet
28 160, but instead is to begin playing the track frames for which begin with the next packet, that is,

1 packet 162. As with Packet Sequence A, in Packet Sequence B the audio information channel device
2 23, in packet 162 and 163, multi-casts frames whose time stamps indicate that they are to be played
3 when the frames that were multi-cast in packets 160 and 161 were to be played. By use of the
4 resynchronize command and specifying a packet in this manner, the audio information channel
5 device can cancel playback of a track for which playback has not yet begun.

6 It will be appreciated that the resynchronize command is generally not necessary for
7 cancelling play back of a track that the audio information channel device 23 has not started multi-
8 casting to the synchrony group 20, since the audio information channel device 23 itself can re-order
9 the play list to accommodate the cancellation.

10 The invention provides a number of advantages. In particular, the invention provides a
11 network audio system in which a number of devices share information can reproduce audio
12 information synchronously, notwithstanding the fact that packets, which may contain digital audio
13 information, transmitted over the network to the various zone players connected thereto may have
14 differing delays and the zone players operate with independent clocks. Moreover, although the
15 invention has been described in connection with audio information, it will be appreciated that the
16 invention will find utility in connection with any type of isochronous information for which
17 synchrony among devices is desired. The system is such that synchrony groups are created and
18 destroyed dynamically, and in such a manner as to avoid requiring a dedicated device as the master
19 device.

20 It will be appreciated that a number of changes and modifications may be made to the
21 network audio system 10 as described above. For example, although the invention has been
22 described as providing that the audio information channel device 23 provides digital audio
23 information to the members synchrony group 20 that has been encoded using particular types of
24 encoding and compression methodologies, it will be appreciated that the audio information channel
25 device 23 can provide digital audio information to various members of the synchrony group 20 that
26 have been encoded and compressed using different types of encoding and compression
27 methodologies, and, moreover, for which different sampling rates have been used. For example, the

1 audio information channel device 23 may provide digital audio information to the master device 21
 2 and slave devices 22(1) through 22(g_1) using the MP3 methodology at a specified sampling rate, the
 3 digital audio information for the same program to slave devices 22(g_1+1) through 22(g_2) using the
 4 WAV methodology at one specified sampling rate, and to slave devices 22(g_2+1) through 22(G)
 5 using the WAV methodology at another specified sampling rate. In that case, the audio information
 6 channel device 23 can specify the particular encoding and compression methodology that has been
 7 used in the encoding type field 57 associated with each frame and the sampling rate in the sampling
 8 rate field 58. Moreover, since the encoding and compression type and sampling rate are specified
 9 for each frame, the encoding and compression type and sampling rate can be changed from frame
 10 to frame. The audio information channel device 23 may use different multi-cast addresses for the
 11 different encoding and compression types and sampling rates, but it will be appreciated that that
 12 would not be required.

13 It will be appreciated that two advantages of providing that the encoding and compression
 14 methodology and the sampling rate is provided on a frame-by-frame basis, instead of on, for
 15 example, a track-by-track basis, is that that would facilitate a slave device joining the synchrony
 16 group 20 at a frame mid-track, without requiring, for example, the master device 21 or the audio
 17 information channel device 23 to notify it of the encoding and compression methodology or the
 18 sampling rate.

19 Another modification is that, instead of the network communications manager 40 of a
 20 member of a synchrony group 20 generating the updated time stamp T_F^u for a digital audio
 21 information frame by adding the time differential value ΔT to the time stamp T_F associated with a
 22 frame, the network communications manager 40 may instead generate the updated time stamp T_F^u
 23 by subtracting the differential time value ΔT from the member's current time T_s as indicated by the
 24 member's digital to analog converter clock 34 at the time at which the digital audio information is
 25 received. It will be appreciated, however, that there may be variable time delays in processing of
 26 messages by the slave device's network communications manager 40, and so it may be preferable
 27 to generate the time differential value ΔT using the time stamp T_F provided by the audio information
 28 channel device 23.

1 In addition, instead of the network communications manager 40 of a member of a synchrony
2 group generating an updated time stamp to reflect the difference between the times indicated by the
3 member's digital to analog converter clock and the audio information channel device's digital to
4 analog converter clock, the network communications manager 40 can generate the time differential
5 value ΔT and provide it to the member's playback scheduler 32. In that case, the member's network
6 communications manager 40 can store each digital audio information frame along with the time
7 stamp T_F as received from the master device in the audio information buffer 21. The playback
8 scheduler 32 can utilize the time differential value ΔT , and the time stamps T_F associated with the
9 digital audio information frames, to determine when the respective digital audio information frames
10 are to be played. In determining when a digital audio information frame is to be played, the
11 playback scheduler can add the time differential value to the time stamp T_F associated with the
12 digital audio frame, and enable the digital audio frame to be coupled to the digital to analog
13 converter 33 when the time indicated by the sum corresponds to the current time as indicated by the
14 slave device's digital to analog converter clock 34. Alternatively, when the member's digital to
15 analog converter clock 34 updates its current time T_s , the playback scheduler can generate an updated
16 current time T'_s by subtracting the differential time value ΔT from the current time T_s , and using the
17 updated current time T'_s to determine when to play a digital audio information frame.

18 As described above, the members of a synchrony group 20 periodically obtain the audio
19 information channel device's current time value and uses the current time value that it receives from
20 the audio information channel device to periodically update the time differential value ΔT that it uses
21 in updating the time stamps associated with the various frames. It will be appreciated that, if the
22 digital to analog converter clock(s) associated with the member(s) of a synchrony group 20 are
23 ensured to have the same rate as the digital to analog converter clock, a member need only obtain
24 the current time value from the audio information channel device once, at the beginning of playback.

25 As another alternative, if the zone players are provided with digital to analog converter clock
26 34 whose time and rate can be set by an element such as the network communications manager 40,
27 when a zone player 11(n) is operating as a member of a synchrony group 20, its network
28 communications manager 40 can use the various types of timing information that it receives from

1 the audio information channel device 23, including the current time information and the playback
 2 timing information indicated by the time stamps that are associated with the various frames 51(f)
 3 comprising the audio and playback timing information that it receives, to adjust the synchrony group
 4 member's digital to analog converter clock's time value and/or the clock rate that it uses for playback.
 5 If the clock's time value is to be adjusted, when the synchrony group member's network
 6 communications manager 40 initially receives the current time information from the audio
 7 information channel device 23 for the synchrony group 20, the network communications manager
 8 40 can set the synchrony group member's digital to analog converter clock 34 to the current time
 9 value as indicated by the audio information channel device's current time information. The network
 10 communications manager 40 can set the clock 34 to the current time value indicated by the audio
 11 information channel device's current time information once, or periodically as it receives the current
 12 time information.

13 Alternatively or in addition, the synchrony group member's network communications
 14 manager 40 can use one or both of the current time information and/or the playback timing
 15 information in the time stamps associated with the respective frames 51(f) to adjust the clock rate
 16 of the clock 34 that it uses for playback. For example, when the synchrony group member's network
 17 communications manager 40 receives a frame 51(f_x) having a time stamp having a time value T_{f_x} ,
 18 it can generate the updated time value $T_{f_x}^U = T_{f_x} + \Delta T$ as described above, and store the frame
 19 with the time stamp with the updated time value in the audio information buffer 30. In addition,
 20 since both the number of samples in a frame and the sampling rate, which determines the rate at
 21 which the frame is to be played, are known to the network communications manager 40, it can use
 22 that information, along with the updated time value $T_{f_x}^U$ that is to be used for frame 51(f_x) to
 23 generate an expected updated time value $T_{f_{x+1}}^E$ that is expected for the updated time stamp of the
 24 next frame 51(f_{x+1}). After the synchrony group member's network communications manager 40
 25 receives the next frame 51(f_{x+1}), it can generate the updated time value $T_{f_{x+1}}^U$ and compare that

1 value to the expected updated time value $T^E_{f_{x+1}}$. If the two time values do not correspond, or if
 2 the difference between them is above a selected threshold level, the clock that is used by the audio
 3 information channel device 23 to generate the time stamps is advancing at a different rate than the
 4 synchrony group member's digital to analog converter clock 34, and so the network communications
 5 manager 40 can adjust the rate of the digital to analog converter clock 34 to approach that of the
 6 clock used by the audio information channel device 23 so that the differential time value ΔT is
 7 constant. On the other hand, if the two time values do correspond, then the time differential value
 8 ΔT is constant, or the difference is below a threshold level, and the network communications
 9 manager 40 need not change the clock rate of the digital to analog converter clock 34. It will be
 10 appreciated that, if the clock rate is to be adjusted, the rate adjustment can be fixed, or it can vary
 11 based on, for example, the difference between the updated time value $T^U_{f_{x+1}}$ and the expected
 12 updated time value $T^E_{f_{x+1}}$.

13 It will also be appreciated that, if no rate adjustment is performed for one frame 51(f_{x+1}), the
 14 synchrony group member's network communications manager 40 can generate an expected updated
 15 time value $T^E_{f_{x+2}}$ that is expected for the updated time stamp of the next frame 51(f_{x+2}) using the
 16 updated time value $T^U_{f_x}$ determined for frame 51(f_x), along with the number of samples in a frame
 17 and the sampling rate, and compare the expected updated time value $T^E_{f_{x+2}}$ to the updated time
 18 value $T^U_{f_{x+2}}$ that it generates when it receives frame 51(f_{x+2}). At that point, if the network
 19 communications manager 41 determines that two time values do not correspond, or if the difference
 20 between them is above a selected threshold level, it can adjust the rate of the digital to analog
 21 converter clock 34. Similar operations can be performed if no rate adjustment is performed for
 22 several successive frames 51(f_{x+1}), 51(f_{x+2}),.... This will accommodate the possibility that the rate
 23 differential between the clock 34 and the clock used by the audio information channel device 23 in

1 generating the time stamps have rates that differ by an amount sufficiently small that it cannot be
2 detected using time stamps of two or more successive frames.

3 Instead or in addition to adjusting the clock rate as described above, the synchrony group
4 member's network communications manager 40 can perform similar operations in connection with
5 adjusting the clock rate in connection with the current time information that it receives from the
6 audio information channel device 23.

7 Furthermore, although the network audio system 10 has been described such that the master
8 device 21 of a synchrony group 20 can, in response to control information provided thereto by a user
9 through the user interface module 13, provide a notification to a zone player 11(n) that it is to
10 become a member of its synchrony group 20 as a slave device 22(g), it will be appreciated that the
11 user interface module 13 can provide the notification directly to the zone player 11(n) that is to
12 become a member of the synchrony group 20. In that case, the zone player 11(n) can notify the
13 master device 21 that it is to become a slave device 22(g) in the synchrony group 20, after which the
14 master device 21 can provide information regarding the synchrony group 20, including the multi-cast
15 and unicast addresses of the audio information channel device and other information as described
16 above.

17 Similarly, although the network audio system 10 has been described such that the master
18 device 21 of a synchrony group 20 can, in response to control information provided thereto by a user
19 through the user interface module 13, provide a command to a slave device 22(g) to enable the slave
20 device 22(g) to adjust its volume, it will be appreciated that the user interface module 13 can provide
21 control information directly to the slave device 22(g) to enable the slave device 22(g) to adjust its
22 volume.

23 In addition, although the network audio system 10 has been described such that each frames
24 51(f) is associated with a frame sequence number (reference field 56, FIG. 4), it will be appreciated
25 that, if the packets described above in connection with Packet Sequence A and Packet Sequence B
26 are provided with packet sequence numbers, the frame sequence numbers need not be provided,
27 since the packet sequence numbers can suffice for defining the frame sequencing.

1 Furthermore, although the network audio system 10 has been described such that the zone
2 players 11(n) are provided with an audio amplifier 35 for amplifying the analog signal provided by
3 the respective digital to analog converters 33, it will be appreciated that a zone player may be
4 provided that does not itself include an audio amplifier. In that case, the analog signal may be
5 coupled to an external amplifier for amplification as necessary before being provided to the audio
6 reproduction device(s) 15(n)(r). It will be appreciated that a single zone player 11(n) may be
7 provided with multiple audio amplifiers and audio reproduction device interfaces, and, if necessary,
8 multiple digital to analog converters 33, to provide audio programs for corresponding numbers of
9 synchrony groups.

10 Similarly, although the zone players 11(n) have been described such that they may be
11 connected to one or more audio information sources, it will be appreciated that an audio information
12 source may form part of and be integrated into a zone player 11(n). For example, a zone player may
13 include a compact disk player, cassette tape player, broadcast radio receiver, or the like, that has been
14 integrated into it. In addition, as noted above, an individual zone player 11(n) may be connected to
15 multiple audio information sources and may contemporaneously operate as the audio information
16 channel device 23 for multiple synchrony groups .

17 In addition, although FIG. 1 shows the network audio system 10 as including one user
18 interface module 13, it will be appreciated that the system 10 may include a plurality of user
19 interface modules. Each user interface module be useful for controlling all of the zone players as
20 described above, or alternatively one or more of the user interface modules may be useful for
21 controlling selected subsets of the zone players.

22 Moreover, it will be appreciated that, although the invention has been described in connection
23 with audio information, it will be appreciated that the invention will find utility in connection with
24 any type of information for which synchrony among devices connected to a network is desired.

25 As noted above, while a zone player 11(n) is operating as audio information channel device
26 23 for a synchrony group 20, when the zone player 11(n)'s audio information source interface 30 or
27 network communications manager 40 stores digital audio information frames based on audio

1 information from an audio information source 14(n)(s) in the audio information buffer 31, it will
2 provide time stamps for the respective frames to schedule them for playback after some time delay
3 after they have been buffered in the audio information buffer 31. The delay is provided so that, for
4 other zone players 11(n'), 11(n''),... that are operating as members of a synchrony group, there will
5 be sufficient time for the audio and playback timing information to be transferred over the network
6 12 to those other zone players 11(n'), 11(n''),... so that it can be processed and played by them at the
7 appropriate time as described above. The time period that is selected for the time delay may be fixed
8 or variable, and in either case may be based on a number of factors. If the time period selected for
9 the time delay is fixed, it may be based on, for example, factors such as an estimate of the maximum
10 latency in the network 12, the estimated maximum loading of the various components comprising
11 the zone players 11(n), and other estimates as will be appreciated by those skilled in the art.

12 The time delay may be the same for audio information from all types of audio information
13 sources, and may be constant over the entire period that the synchrony group 20 is playing an audio
14 work. Alternatively, different time delays may be utilized based on various criteria. For example,
15 if the audio information is to be played independently of information associated with other types of
16 media, the time delay may be selected to be relatively long, on the order of a significant fraction of
17 a second, or longer. On the other hand, if the audio information is to be played contemporaneously
18 with, for example, video information, which may be supplied by, for example, a video disk, video
19 tape cassette, over cable, satellite, or broadcast television, which may not be buffered or which may
20 be displayed independently of the network audio system 10, it may be undesirable to provide for
21 such a lengthy delay, since the time delay of the audio playback, in relation to the video display, may
22 be noticeable. In that case, the zone player 11(n) may provide for a much shorter time delay. In one
23 embodiment, the time delay provided for audio information to be played concurrently with video
24 information is selected to be generally on the order of fifty milliseconds, which would barely, if at
25 all, be perceptible to someone viewing the video. Other desirable time delays for information from
26 other types of sources will be apparent to those skilled in the art.

27 As yet a further possibility, the zone player 11(n), when operating as an audio information
28 channel device 23 for a synchrony group 20, can dynamically determine the time delay based on a

1 number of conditions in the network audio system 10, including, for example, the message transfer
2 latency in network 12, the loading of microprocessors or other components that are used in the
3 various zone players 11(n'), 11(n''),... that may comprise a synchrony group 20, as well as other
4 factors. For example, if the audio information channel device 23 determines that the latency in the
5 network 12 has increased beyond a selected threshold, the audio information channel device 23 can
6 adjust the delay to increase the likelihood that the members of the synchrony group 20 will be able
7 to receive the packets and process the frames so that they will be able to play them at the appropriate
8 times. Similarly, if the audio information channel device 23 is notified that a member of the
9 synchrony group 20 to which it provides audio information requires additional time to receive and
10 process the frames that it transmits, the audio information channel device 23 can adjust the delay
11 accordingly. It will be appreciated that, to reduce or minimize possible discontinuities in the audio
12 playback by the members of the synchrony group, the audio information channel device 23 can,
13 instead of adjusting the time delay during a particular audio track, adjust the time delay between
14 tracks, during silent periods of a track or otherwise as will be appreciated by those skilled in the art.
15 In addition, the audio information channel device 23 can use conventional audio compression
16 methodologies to facilitate a speeding up and/or slowing down of playback of an audio track while
17 it is in the process of providing additional time delay. Generally, the members of the synchrony
18 group 20 can provide notifications to the audio information channel device 23 if they determine that
19 they will need an additional time delay, and the audio information channel device 23 can adjust the
20 time delay in accordance with the notifications from the members of the synchrony group 20.

21 It will be appreciated that a system in accordance with the invention can be constructed in
22 whole or in part from special purpose hardware or a general purpose computer system, or any
23 combination thereof, any portion of which may be controlled by a suitable program. Any program
24 may in whole or in part comprise part of or be stored on the system in a conventional manner, or it
25 may in whole or in part be provided in to the system over a network or other mechanism for
26 transferring information in a conventional manner. In addition, it will be appreciated that the system
27 may be operated and/or otherwise controlled by means of information provided by an operator using
28 operator input elements (not shown) which may be connected directly to the system or which may

1 transfer the information to the system over a network or other mechanism for transferring
2 information in a conventional manner.

3 The foregoing description has been limited to a specific embodiment of this invention. It will
4 be apparent, however, that various variations and modifications may be made to the invention, with
5 the attainment of some or all of the advantages of the invention. It is the object of the appended
6 claims to cover these and such other variations and modifications as come within the true spirit and
7 scope of the invention.

8 What is claimed as new and desired to be secured by Letters Patent of the United States is: